

Installation  
and  
Configuration  
Guide

# HP StorageWorks HSG80 ACS Solution Software V8.8

## for Tru64 UNIX

**Product Version:** 8.8-1

First Edition (March 2005)

**Part Number:** AA-RV1VA-TE

This guide provides installation and configuration instructions and reference material for operation of the HSG80 ACS Solution Software V8.8-1 for Tru64 UNIX.



© Copyright 2000-2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Compaq Computer Corporation is a wholly-owned subsidiary of Hewlett-Packard Company.

Intel® is a U.S. registered trademark of Intel Corporation.

Microsoft®, MS-DOS®, MS Windows®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

**HSG80 ACS Solution Software V8.8 for Tru64 UNIX Installation and Configuration Guide**

**First Edition (March 2005)**

**Part Number: AA-RV1VA-TE**

## contents

<b>About this Guide</b> .....	<b>13</b>
Overview .....	14
Intended Audience .....	14
Related Documentation .....	14
Chapter Content Summary .....	16
Conventions .....	18
Document conventions .....	18
Symbols in Text .....	18
Symbols on Equipment .....	19
Rack Stability .....	20
Getting Help .....	21
Technical Support .....	21
Storage web site .....	21
HP Authorized Reseller .....	22
Configuration Flowchart .....	22
<b>1 Planning a Subsystem</b> .....	<b>27</b>
Defining Subsystems .....	28
Controller Designations A and B .....	28
Controller Designations “This Controller” and “Other Controller” .....	29
What is Failover Mode? .....	31
Transparent Failover Mode .....	31
Multiple-Bus Failover Mode .....	33
Selecting a Cache Mode .....	36
Read Caching .....	36
Read-Ahead Caching .....	36
Write-Back Caching .....	36
Write-Through Caching .....	36
Enabling Mirrored Caching .....	37
What is the Command Console LUN? .....	38

Determining the Address of the CCL.....	38
Enabling/Disabling the CCL in SCSI-2 Mode.....	38
Enabling the CCL .....	38
Disabling the CCL.....	39
Enabling/Disabling CCL in SCSI-3 Mode.....	39
Determining Connections.....	39
Naming Connections .....	39
Numbers of Connections .....	40
Assigning Unit Numbers .....	44
Matching Units to Host Connections in Transparent Failover Mode .....	44
Matching Units to Host Connections in Multiple-Bus Failover Mode .....	46
Assigning Unit Numbers Depending on SCSI_VERSION .....	47
Assigning Host Connection Offsets and Unit Numbers in SCSI-3 Mode.....	47
Assigning Host Connection Offsets and Unit Numbers in SCSI-2 Mode.....	47
What is Selective Storage Presentation?.....	48
Restricting Host Access in Transparent Failover Mode .....	49
Restricting Host Access by Separate Links.....	49
Restricting Host Access by Disabling Access Paths .....	50
Restricting Host Access by Offsets .....	51
Restricting Host Access in Multiple-Bus Failover Mode .....	52
Enabling the Access Path of Selected Host Connections .....	52
Restricting Host Access by Offsets .....	54
Worldwide Names (Node IDs and Port IDs).....	55
Restoring Worldwide Names (Node IDs) .....	56
Unit Worldwide Names (LUN IDs).....	57
<b>2 Planning Storage Configurations .....</b>	<b>59</b>
Where to Start.....	60
Determining Storage Requirements .....	61
Configuration Rules for the Controller.....	61
Addressing Conventions for Device PTL.....	62
Examples - Model 2200 Storage Maps, PTL Addressing .....	65
Choosing a Container Type .....	72
Creating a Storageset Profile .....	75
Planning Considerations for Storageset .....	76
Stripeset Planning Considerations .....	76
Mirrorset Planning Considerations .....	80
Keep these points in mind when planning mirrorsets .....	81
RAIDset Planning Considerations .....	81

Keep these points in mind when planning RAIDsets .....	82
Striped Mirrorset Planning Considerations .....	83
Storageset Expansion Considerations .....	85
Partition Planning Considerations .....	85
Defining a Partition .....	86
Guidelines for Partitioning Storagesets and Disk Drives .....	86
Changing Characteristics Through Switches .....	87
Enabling Switches .....	87
Changing Switches .....	87
Specifying Storageset and Partition Switches .....	88
RAIDset Switches .....	88
Mirrorset Switches .....	88
Partition Switches .....	88
Specifying Initialization Switches .....	89
Chunk Size .....	89
Increasing the Request Rate .....	90
Increasing Sequential Data Transfer Performance .....	91
Save Configuration .....	91
Destroy/NodeDestroy .....	92
Geometry .....	93
Specifying Unit Switches .....	93
Creating Storage Maps .....	93
Using LOCATE Command to Find Devices .....	94
Example Storage Map—Model 4310R Disk Enclosure .....	95
<b>3 Preparing the Host System .....</b>	<b>99</b>
Installing RAID Array Storage System .....	100
Making a Physical Connection .....	104
Preparing to Install Host Bus Adapter .....	104
Installing Host Bus Adapter .....	104
Verifying Disk Array Device Recognition .....	105
Example of calculation .....	105
Configure Disk Devices .....	106
Writing the Partition Labels .....	106
Creating the File Systems .....	107
Creating a File System For Each Device .....	107
Creating an Advanced File System .....	107
To create a new domain: .....	107
Creating Mount Directories .....	108

Mounting the File Systems . . . . .	108
To mount a Tru64 file system . . . . .	108
To mount an advanced file system . . . . .	109
Verifying the File Systems . . . . .	109
Setting and Verifying Automatic Mounting at Bootup . . . . .	110
Checking path failover . . . . .	111
DECsafe Available Server Environment (ASE) . . . . .	112
Using genvmunix . . . . .	112
HSG80 Units and Tru64 UNIX Utilities . . . . .	112
Add recognition for Fibre Channel adapters . . . . .	112
Solution Software Upgrade Procedures . . . . .	113
Installing Agent . . . . .	113
Upgrading the Agent and ACS on Standalone Servers . . . . .	114
New Features, ACS 8.8-1 for Tru64 . . . . .	115
Host Connection Table Management Improvements . . . . .	116
Host Connection Table Locking . . . . .	116
Viewing Host Connection Table Lock State . . . . .	116
Adding Rejected Host Connections to Locked Host Connection Table . . . . .	119
Implementation Notes . . . . .	119
Selective Management Presentation . . . . .	120
Removing Management Agent Host systems . . . . .	120
Adding Management Agent Host systems . . . . .	121
Display Enabled Management Agents . . . . .	121
Enabling SAN Security . . . . .	122
Linking WWIDs for Snap and Clone Units . . . . .	123
CLI format . . . . .	123
Implementation Notes . . . . .	124
SMART Error Eject . . . . .	125
Error Threshold for Drives . . . . .	129
<b>4 Installing and Configuring the HSG Agent . . . . .</b>	<b>131</b>
Why Use StorageWorks Command Console (SWCC)? . . . . .	132
Installation and Configuration Overview . . . . .	133
About the Network Connection for the Agent . . . . .	134
Before Installing the Agent . . . . .	136
Downloading the Host Kit Software From the Web . . . . .	137
Installing and Configuring the Agent . . . . .	137
Configure Client System Options . . . . .	138
Enter Subsystem Information . . . . .	140

---

Configure Email Notification Details .....	140
Final System Configuration.....	141
Reconfiguring the Agent.....	143
Removing the Agent.....	145
<b>5 FC Configuration Procedures .....</b>	<b>147</b>
Establishing a Local Connection .....	148
Setting Up a Single Controller .....	149
Powering On and Establishing Communication .....	149
Cabling a Single Controller .....	149
Configuring a Single Controller Using CLI.....	150
Verifying the Node ID and Checking for Any Previous Connections .....	150
Configuring Controller Settings .....	151
Restarting the Controller .....	152
Setting Time and Verifying all Commands.....	152
Plugging in the FC Cable and Verify Connections.....	155
Repeating Procedure for Each Host Adapter.....	155
Verifying Installation.....	155
For V5.1x Use .....	155
Setting Up a Controller Pair .....	156
Powering Up and Establishing Communication .....	156
Cabling a Controller Pair .....	156
Configuring a Controller Pair Using CLI.....	157
Configuring Controller Settings .....	158
Restarting the Controller .....	159
Setting Time and Verifying All Commands .....	159
Plugging in the FC Cable and Verify Connections .....	162
Repeating Procedure for Each Host Adapter Connection .....	162
Verifying Installation.....	162
For V5.1x Use .....	163
Configuring Devices.....	163
Configuring Storage Containers .....	163
Configuring a Stripeset .....	164
Configuring a Mirrorset .....	165
Configuring a RAIDset .....	165
Configuring a Striped Mirrorset .....	166
Configuring a Single-Disk Unit (JBOD) .....	167
Configuring a Partition.....	167
Assigning Unit Numbers and Unit Qualifiers .....	169

Assigning a Unit Number to a Storageset . . . . .	169
Assigning a Unit Number to a Single (JBOD) Disk . . . . .	169
Assigning a Unit Number to a Partition . . . . .	169
Preferring Units . . . . .	170
Configuration Options . . . . .	170
Changing the CLI Prompt . . . . .	170
Mirroring cache . . . . .	170
Adding Disk Drives . . . . .	171
Adding a Disk Drive to the Spareset . . . . .	171
Removing a Disk Drive from the Spareset . . . . .	171
Enabling Autospare . . . . .	172
Deleting a Storageset . . . . .	173
Changing Switches for a Storageset or Device . . . . .	173
Displaying the Current Switches . . . . .	173
Changing RAIDset and Mirrorset Switches . . . . .	173
Changing Device Switches . . . . .	174
Changing Initialize Switches . . . . .	174
Changing Unit Switches . . . . .	174
Verifying Storage Configuration from Host . . . . .	175
<b>6 Using CLI for Configuration . . . . .</b>	<b>177</b>
CLI Configuration Example . . . . .	181
<b>7 Backing Up, Cloning, and Moving Data . . . . .</b>	<b>187</b>
Backing Up Subsystem Configurations . . . . .	188
Creating Clones for Backup . . . . .	189
Moving Storagesets . . . . .	193
<b>A Subsystem Profile Templates . . . . .</b>	<b>197</b>
Storageset Profile . . . . .	198
Storage Map Template 1 for the BA370 Enclosure . . . . .	200
Storage Map Template 2 for the Second BA370 Enclosure . . . . .	201
Storage Map Template 3 for the TSingle-Bushird BA370 Enclosure . . . . .	202
Storage Map Template 4 for the Model 4214R Disk Enclosure . . . . .	203
Storage Map Template 5 for the Model 4254 Disk Enclosure . . . . .	205
Storage Map Template 6 for the Model 4310R Disk Enclosure . . . . .	207
Storage Map Template 7 for the Model 4350R Disk Enclosure . . . . .	209
Storage Map Template 8 for the Model 4314R Disk Enclosure . . . . .	210
Storage Map Template 9 for the Model 4354R Disk Enclosure . . . . .	212

<b>B</b>	<b>Installing, Configuring, and Removing the Client</b>	<b>213</b>
	Why Install the Client	214
	Before You Install the Client	215
	Installing the Client	216
	Installing the Integration Patch	217
	Should I Install the Integration Patch?	217
	Performing the Integration Patch Installation	217
	Integrating Controller's SWCC Storage Window with CIM	218
	Finding the Controller's Storage Window	219
	Troubleshooting Client Installation	219
	Invalid Network Port Assignments During Installation	219
	“There is no disk in the drive” Message	220
	Adding Storage Subsystem and its Host to Navigation Tree	221
	Removing Command Console Client	223
	Where to Find Additional Information	224
	About the User Guide	224
	About the Online Help	224
<b>C</b>	<b>SWCC Agent in TruCluster Environment</b>	<b>225</b>
	SWCC Overview	226
	Running the SWCC Agent under ASE Services (UNIX 4.0 only)	227
	Creating the Start/Stop Script	228
	Use the ASEMGR Utility to Create the Service	229
	Selecting an Automatic Service Placement (ASP) Policy	235
	Running the SWCC Agent on a V5.x Cluster	237
	Problems with Running the Agent on Multiple Clusters	240
	Configure the Controller	241
	Use Multiple-Bus Failover Mode	241
	Verify That the HSG80/HSG60 Unit Offsets Are Zero	243
	Use SCSI-3 Mode	244
	Install and Run the Agent on One Cluster Member	245
	Example of Installing the Agent on a Cluster Member	248
	Create the CAA Action Script	254
	Complete CAA Action Script	255
	Create the CAA Resource	261
	Register the CAA Resource	262
	Example of Creating, Registering, and Starting the CAA Resource	263
	Viewing Events Posted by the Action Script	264
	Edit the Startup Script	265

Install and Configure the SWCC Client .....	265
Do Not Use <code>swcc_config</code> to Start, Stop, or Restart .....	266
Verifying Success .....	267
<b>Glossary</b> .....	<b>269</b>
<b>Index</b> .....	<b>291</b>
<b>Figures</b>	
1 General configuration flowchart (panel 1) .....	23
2 General configuration flowchart (panel 2) .....	24
3 SWCC storage configuration flowchart (panel 3) .....	25
4 Location of controllers and cache modules in a Model 2200 enclosure .....	28
5 Location of controllers and cache modules in a BA370 enclosure .....	29
6 “This controller” and “other controller” for the Model 2200 enclosure .....	30
7 “This controller” and “other controller” for the BA370 enclosure .....	30
8 Transparent failover—normal operation .....	32
9 Transparent failover—after failover from controller B to controller A .....	33
10 Typical multiple-bus configuration .....	35
11 Mirrored caching .....	37
12 Connections in separate-link, transparent failover mode configurations .....	41
13 Connections in single-link, transparent failover mode configurations .....	42
14 Connections in multiple-bus failover mode .....	43
15 LUN presentation to hosts, as determined by offset .....	45
16 Limiting host access in transparent failover mode .....	50
17 Limiting host access in multiple-bus failover mode .....	53
18 Placement of the worldwide name label on the Model 2200 enclosure .....	56
19 Placement of the worldwide name label on the BA370 enclosure .....	57
20 Mapping a unit to physical disk drives .....	63
21 PTL naming convention .....	63
22 How data is laid out on disks in BA370 enclosure configuration .....	64
23 Storage container types .....	73
24 Three-member RAID 0 stripeset (example 1) .....	77
25 Three-member RAID 0 stripeset (example 2) .....	78
26 Mirrorsets maintain two copies of the same data (example 1) .....	80
27 Mirrorset (example 2) .....	80
28 Five-member RAIDset using parity .....	82
29 Striped mirrorset (example 1) .....	84

30	Striped mirrorset (example 2) . . . . .	84
31	One example of a partitioned single-disk unit . . . . .	85
32	Large chunk size increases request rate . . . . .	90
33	Dual-Bus Enterprise Storage RAID Array Storage System. . . . .	102
34	Single-Bus Enterprise Storage RAID Array Storage System . . . . .	103
35	An example of a network connection. . . . .	135
36	Maintenance port connection . . . . .	148
37	Single controller cabling . . . . .	150
38	Controller pair failover cabling . . . . .	157
39	Storage container types . . . . .	164
40	Example storage map for the BA370 Enclosure . . . . .	178
41	Example, three non-clustered host systems . . . . .	179
42	Representative multiple-bus failover . . . . .	180
43	Example, logical or virtual disks comprised of storagesets. . . . .	181
44	CLONE utility steps for duplicating unit members . . . . .	190
45	Navigation Window . . . . .	221
46	Navigation window showing storage host system “Atlanta” . . . . .	222
47	Navigation window showing expanded “Atlanta” host icon . . . . .	222

## Tables

1	Summary of Chapter Contents. . . . .	16
2	Document Conventions . . . . .	18
3	Unit Assignments and SCSI_VERSION . . . . .	48
4	PTL addressing, single-bus configuration, six Model 4320R enclosures. . . . .	67
5	PTL addressing, dual-bus configuration, three Model 4350R enclosures . . . . .	69
6	PTL addressing, single-bus configuration, six Model 4314R enclosures . . . . .	70
7	PTL addressing, dual-bus configuration, three Model 4354A enclosures . . . . .	71
8	Comparison of Container Types . . . . .	74
9	Example of Storageset Profile . . . . .	75
10	Example Chunk Sizes . . . . .	91
11	Model 4310 disk enclosure, example of storage map . . . . .	95
12	Boot time mounting parameters. . . . .	110
13	Mounting parameter descriptions. . . . .	110
14	SWCC Features and Components . . . . .	132
15	Installation and Configuration Overview. . . . .	133
16	Client System Access Options . . . . .	139
17	Client System Notification Options . . . . .	139

18	Definitions of Email Notification Options . . . . .	141
19	Information Needed to Configure Agent . . . . .	144
20	Configuration Requirements . . . . .	237
21	Required Configuration Details . . . . .	245
22	Troubleshooting . . . . .	268

## about this guide

This Installation Guide describes how to install and configure the HSG80 ACS Solution Software Version 8.8-1 for Tru64 UNIX.

This guide provides information to help you:

- Plan the storage array subsystem
- Install and configure the storage array subsystem on individual operating system platforms

“About this Guide” topics include:

- [Overview](#), page 14
- [Chapter Content Summary](#), page 16
- [Conventions](#), page 18
- [Rack Stability](#), page 20
- [Getting Help](#), page 21
- [HP Authorized Reseller](#), page 22
- [Configuration Flowchart](#), page 22

This book does not contain information about the operating environments to which the controller may be connected, nor does it contain detailed information about subsystem enclosures or their components. See the documentation that accompanied these peripherals for information about them.

## Overview

This section covers the following topics:

[Intended Audience](#), page 14

[Related Documentation](#), page 14

## Intended Audience

This book is intended for use by system administrators and system technicians who have a basic experience with storage and networking.

## Related Documentation

In addition to this guide, corresponding information can be found in:

- ACS V8.8-1 controller documentation (software delivered via PCMCIA cards)
- *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*, EK-G80CL-RA. C01
- *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide*, EK-G80MS-SA. C01
- *HP StorageWorks HSG60 and HSG80 Array Controller Software Troubleshooting Guide*, EK-G80TS-SA. C01
- StorageWorks Command Console (SWCC) v2.5 documentation (client software delivered in solutions kits)
- *HP StorageWorks HP StorageWorks Command Console Version 2.5 User Guide*, AA- RV1UA -TE
- *HP StorageWorks HP StorageWorks Command Console Version 2.5 Release Notes*, AV- RV1TA -TE
- *HP StorageWorks Command Console V2.5 Online Help*, AA-RS20A-TE and AA-RS21A-TE
- Host-specific documentation (SWCC Agent and HBA software delivered in solutions kits)
- Installation and Configuration Guide (platform-specific) - the guide you are reading

- *HP StorageWorks HSG80 ACS Solution Software Release Notes*  
(platform-specific)
- *HP StorageWorks Enterprise/Modular Storage RAID Array Fibre Channel Arbitrated Loop Configurations for Windows, Tru64, and Sun Solaris Application Note* (AA-RS1ZB-TE) - Solution software host support includes the following platforms:
  - IBM AIX
  - HP-UX
  - Linux (Red Hat x86/Alpha, SuSE x86/Alpha, Caldera x86)
  - Novell NetWare
  - Open VMS
  - Sun Solaris
  - Tru64 UNIX
  - Windows NT/2000/Windows Server 2003 (32-bit)

Additional support required by HSG80 ACS Solution Software Version 8.8-1, but delivered through external programs, includes the following:

- Heterogeneous “rules-based” SAN configurations
- Host-Bus Adapter (HBA) products
- Applicable Storage Utility Management Suite (SUMS) components
- Vendor-specific switch products
- Secure Path products (Windows, NetWare, Sun, AIX, HP-UX)
- Data Replication Manager (DRM) under ACS 8.8-1P
- Enterprise Volume Manager (EVM) under ACS 8.8-1S
- Enterprise Backup Solution (EBS)
- Additional ACS variants (W, R)

# Chapter Content Summary

[Table 1](#) below summarizes the content of the chapters.

**Table 1: Summary of Chapter Contents**

Chapters	Description
1. Planning a Subsystem	This chapter focuses on technical terms and knowledge needed to plan and implement storage array subsystems.
2. Planning Storage Configurations	Plan the storage configuration of your subsystem, using individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. This chapter describes addressing conventions, configuration rules, creating storage profiles, and creating storage maps.
3. Preparing the host system	How to prepare your Tru64 UNIX host computer to accommodate the HSG80 controller storage subsystem.
4. Installing and Configuring the HSG Agent	The Agent for HSG for a specific operating system polls the storage.
5. FC Configuration Procedures	How to configure a subsystem that uses Fibre Channel (FC) fabric topology. In fabric topology, the controller connects to its hosts through switches. How to verify that multiple paths exist to virtual disk units under Tru64 UNIX V5.1x.
6. Using CLI for Configuration	How-to example of configuring a storage subsystem using the Command Line Interpreter (CLI).
7. Backing Up, Cloning, and Moving Data	Description of common procedures that are not mentioned elsewhere in this guide. Backing Up Subsystem Configuration Cloning Data for Backup Moving Storagesets

**Table 1: Summary of Chapter Contents (Continued)**

Chapters	Description
Appendix A. Subsystem Profile Templates	This appendix contains storage set profiles to copy and use to create your system profiles. It also contains an enclosure template to use to help keep track of the location of devices and storage sets in your shelves. Four (4) templates will be needed for the subsystem.
Appendix B. Installing, Configuring, and Removing the Client	<p>The Client monitors and manages a storage subsystem.</p> <p>The following information is included in this appendix:</p> <ul style="list-style-type: none"> <li>■ Why Install the Client?</li> <li>■ Before You Install the Client</li> <li>■ Installing the Client</li> <li>■ Installing Integration Patch</li> <li>■ Troubleshooting the Client Installation</li> <li>■ Adding Storage Subsystem and its Host to the Navigation Tree</li> <li>■ Removing the Command Console Client</li> <li>■ Where to Find Additional Information</li> </ul>
Appendix C. SWCC in TruCluster Environment	This appendix describes setup of the SWCC Agent on TruClusters Tru64 UNIX Version 5.X.

## Conventions

Conventions consist of the following:

- [Document conventions](#)
- [Symbols in Text](#)
- [Symbols on Equipment](#)

### Document conventions

This document follows the conventions in [Table 2](#).

**Table 2: Document Conventions**

Element	Convention
Cross-reference links	Blue text: <a href="#">Figure 1</a>
Key names, menu items, buttons, and dialog box titles	<b>Bold</b>
File names, application names, and text emphasis	<i>Italics</i>
User input, command names, system responses (output and messages)	Monospace font COMMAND NAMES are uppercase unless they are case sensitive
Variables	<i>Monospace, italic font</i>
web site addresses	Sans serif font ( <a href="http://www.hp.com">http://www.hp.com</a> )

### Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.



**WARNING:** Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.

---



**Caution:** Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

---

---

**Note:** Text set off in this manner presents commentary, sidelights, or interesting points of information.

---

## Symbols on Equipment

---



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

**WARNING:** To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

---



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

**WARNING:** To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

---



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

**WARNING:** To reduce the risk of injury from a hot component, allow the surface to cool before touching.

---



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

**WARNING:** To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

---



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

**WARNING:** To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

---

## Rack Stability



**WARNING:** To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- In single rack installations, the stabilizing feet are attached to the rack.
- In multiple rack installations, the racks are coupled.
- Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.

---

## Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site.

## Technical Support

Telephone numbers for worldwide technical support are listed on the following HP web site: <http://www.hp.com/support/>. From this web site, select the country of origin.

---

**Note:** For continuous quality improvement, calls may be recorded or monitored.

---

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the following web site: <http://www.hp.com>.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

## Storage web site

The Storage web site has the latest information on this product, as well as the latest drivers. Access the Storage web site at: <http://www.hp.com/storage>. From this web site, select the appropriate product or solution.

## HP Authorized Reseller

For the name of your nearest authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the Storage web site for locations and telephone numbers

## Configuration Flowchart

A three-part flowchart ([Figure 1](#), [Figure 2](#), and [Figure 3](#)) is shown on the following pages. Refer to these charts while installing and configuring a new storage subsystem. All references in the flowcharts pertain to pages in this guide, unless otherwise indicated.

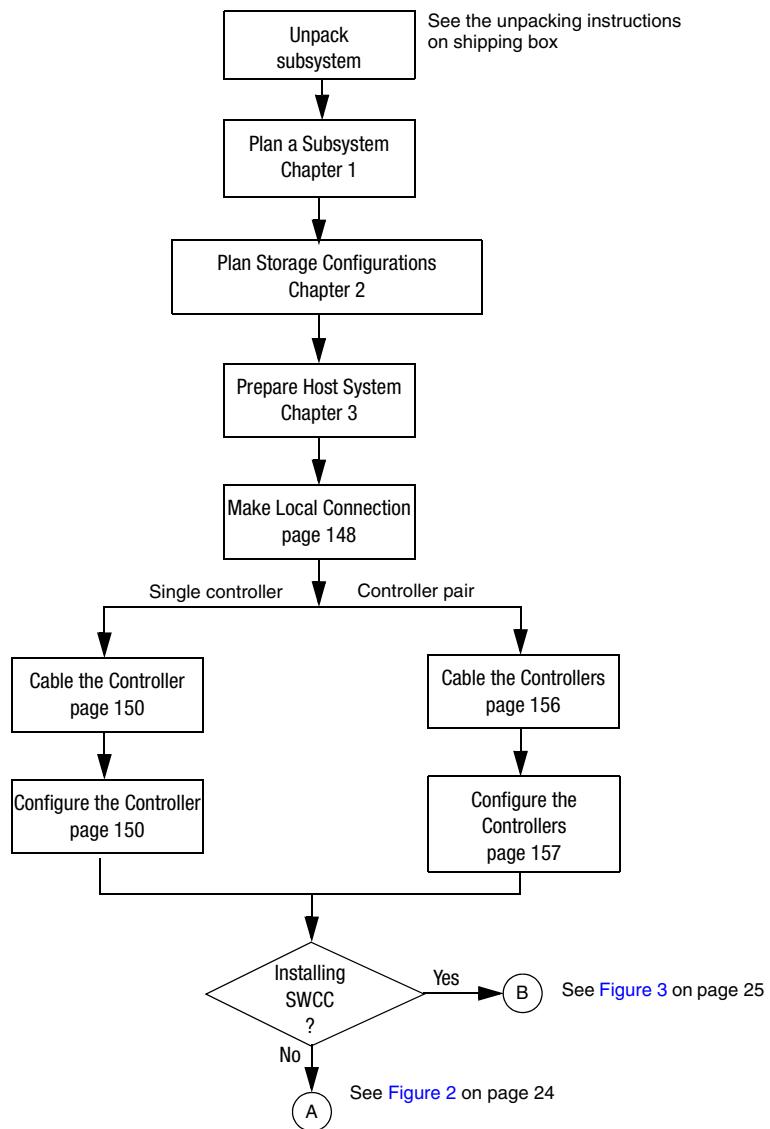
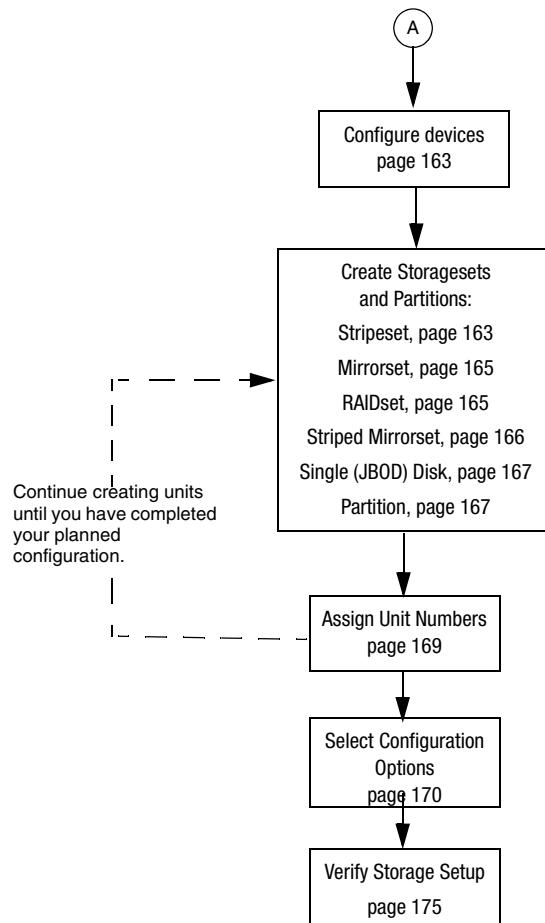
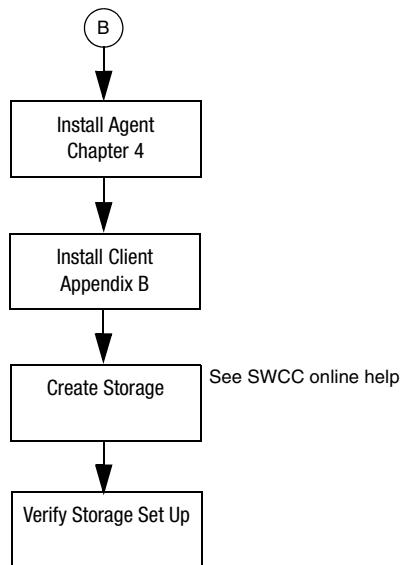


Figure 1: General configuration flowchart (panel 1)



**Figure 2: General configuration flowchart (panel 2)**



**Figure 3: SWCC storage configuration flowchart (panel 3)**



# Planning a Subsystem

This chapter provides information that helps you plan how to configure the storage array subsystem. This chapter focuses on the technical terms and knowledge needed to plan and implement storage subsystems.

---

**Note:** This chapter frequently references the command line interface (CLI). For the complete syntax and descriptions of the CLI commands, see the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

---

The following information is included in this chapter:

- [Defining Subsystems](#), page 28
- [What is Failover Mode?](#), page 31
- [Selecting a Cache Mode](#), page 36
- [Enabling Mirrored Caching](#), page 37
- [What is the Command Console LUN?](#), page 38
- [Determining Connections](#), page 39
- [Assigning Unit Numbers](#), page 44
- [What is Selective Storage Presentation?](#), page 48

---

**Note:** DILX should be run for ten minutes on all units to delete the 8 MB EISA partition. Refer to *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for details.

---

Refer to Chapter 2 when planning the types of storage containers you need. Storage containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives.

## Defining Subsystems

This section describes the terms *this controller* and *other controller*. It also presents graphics of the Model 2200 and BA370 enclosures.

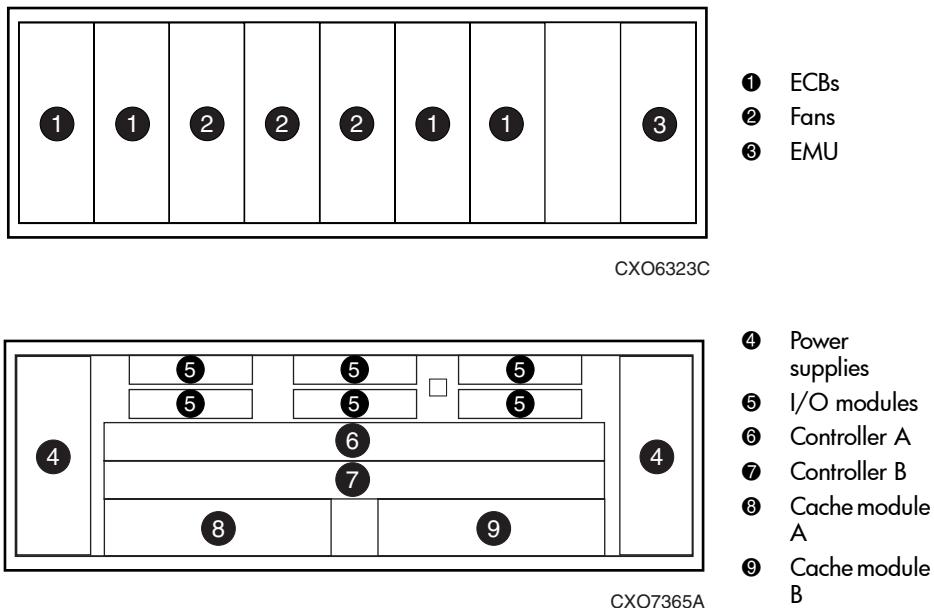
**Note:** The HSG80 controller uses the BA370 or Model 2200 enclosure.

## Controller Designations A and B

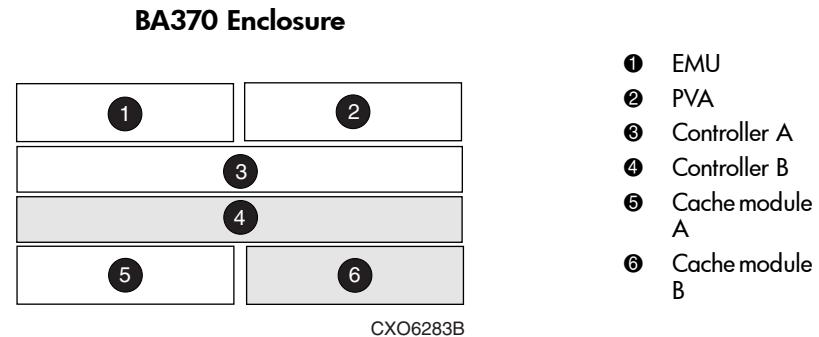
The terms A, B, “this controller,” and “other controller,” respectively are used to distinguish one controller from another in a two-controller (also called dual-redundant) subsystem.

Controllers and cache modules are designated either A or B depending on their location in the enclosure, as shown in [Figure 4](#) for the Model 2200 enclosure and in [Figure 5](#) for the BA370 enclosure.

**Model 2200 Enclosure**



**Figure 4: Location of controllers and cache modules in a Model 2200 enclosure**



**Figure 5: Location of controllers and cache modules in a BA370 enclosure**

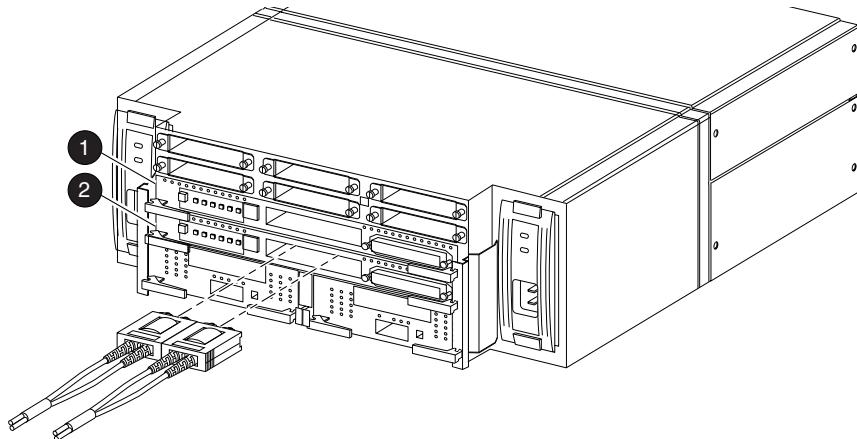
## Controller Designations “This Controller” and “Other Controller”

Some CLI commands use the terms “this” and “other” to identify one controller or the other in a dual-redundant pair. These designations are a shortened form of “this controller” and “other controller.” These terms are defined as follows:

- “**this controller**”—the controller that is the focus of the CLI session. “This controller” is the controller to which the maintenance terminal is attached and through which the CLI commands are being entered. “This controller” can be abbreviated to “this” in CLI commands.
- “**other controller**”—the controller that is not the focus of the CLI session and through which CLI commands are not being entered. “Other controller” can be abbreviated to “other” in CLI commands.

[Figure 6](#) shows the relationship between “this controller” and “other controller” in a Model 2200 enclosure. [Figure 7](#) shows the same relationship in a BA370 enclosure.

**Model 2200 Enclosure**

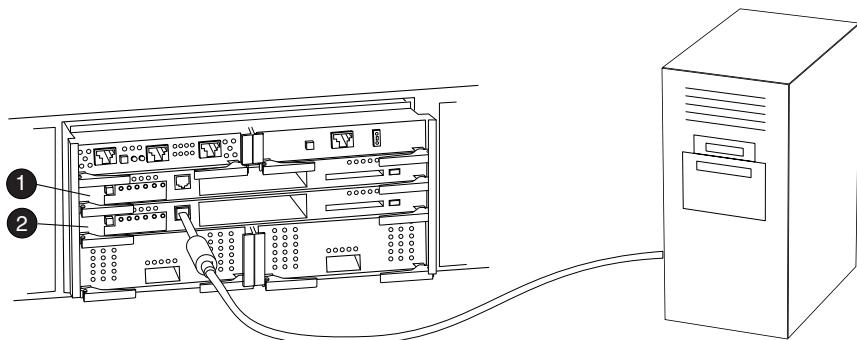


❶ This controller

❷ Other controller

**Figure 6: "This controller" and "other controller" for the Model 2200 enclosure**

**BA370 Enclosure**



❶ Other controller

❷ This controller

**Figure 7: "This controller" and "other controller" for the BA370 enclosure**

## What is Failover Mode?

Failover is a way to keep the storage array available to the host if one of the controllers becomes unresponsive. A controller can become unresponsive because of a controller hardware failure or, in multiple-bus mode only, due to a failure of the link between host and controller or host-bus adapter. Failover keeps the storage array available to the hosts by allowing the surviving controller to take over total control of the subsystem.

There are two failover modes:

- **Transparent**, which is handled by the surviving controller and is invisible (transparent) to the hosts.
- **Multiple-bus**, which is handled by the surviving controller and which is handled by the hosts through either additional software or as a feature of the operating system.

Either mode of failover can work with loop or fabric topology.

### Transparent Failover Mode

Transparent failover mode has the following characteristics:

- Hosts do not know failover has taken place
- Units are divided between host ports 1 and 2

A unit or storage set is a physical or virtual device of the subsystem. It is typically assigned a logical unit number (LUN) and is managed by the HSG80 controller and presented to a server through the Fibre Channel bus and the server's host bus adapter. Disks that are set up as independent disks (JBODs) or RAIDsets are referred to as storage sets.

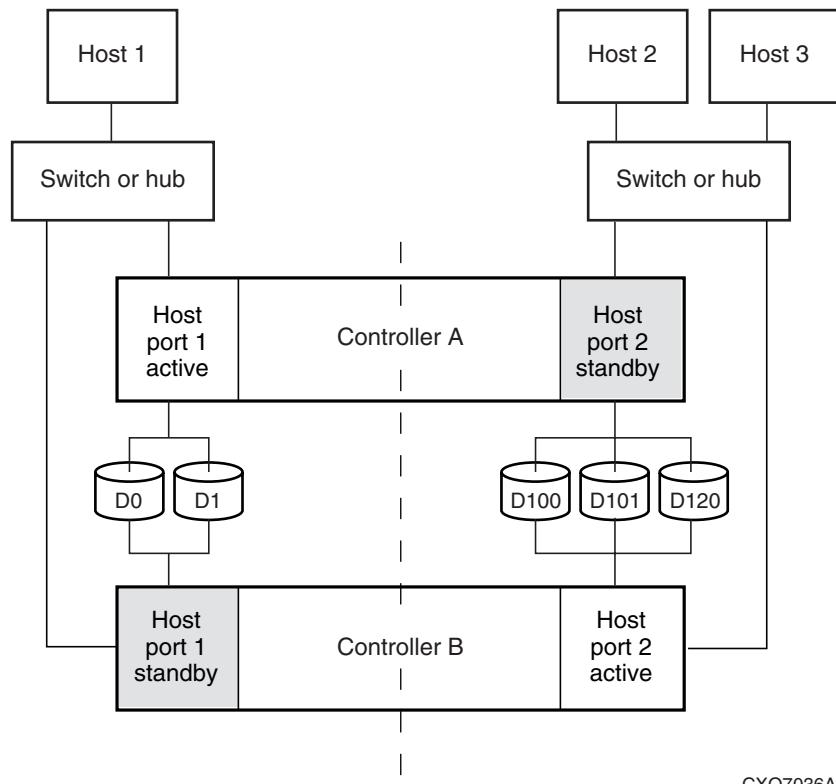
In transparent failover mode, host port 1 of controller A and host port 1 of controller B must be on the same Fibre Channel link. Host port 2 of controller A and host port 2 of controller B must also be on the same Fibre Channel link. Depending on operating system restrictions and requirements, the port 1 link and the port 2 link can be separate links, or they can be the same link.

At any time, host port 1 is active on only one controller, and host port 2 is active on only one controller. The other ports are in standby mode. In normal operation, both host port 1 on controller A and host port 2 on controller B are active. A representative configuration is shown in [Figure 8](#). The active and standby ports share port identity, enabling the standby port to take over for the active one. If one controller fails, its companion controller (known as the surviving controller) takes control by making both its host ports active, as shown in [Figure 9](#).

Units are divided between the host ports:

- Units 0-99 are on host port 1 of both controllers (but accessible only through the active port).
- Units 100-199 are on host port 2 of both controllers (but accessible only through the active port).

Transparent failover only compensates for a controller failure, and not for failures of either the Fibre Channel link or host Fibre Channel adapters.



**Figure 8: Transparent failover—normal operation**

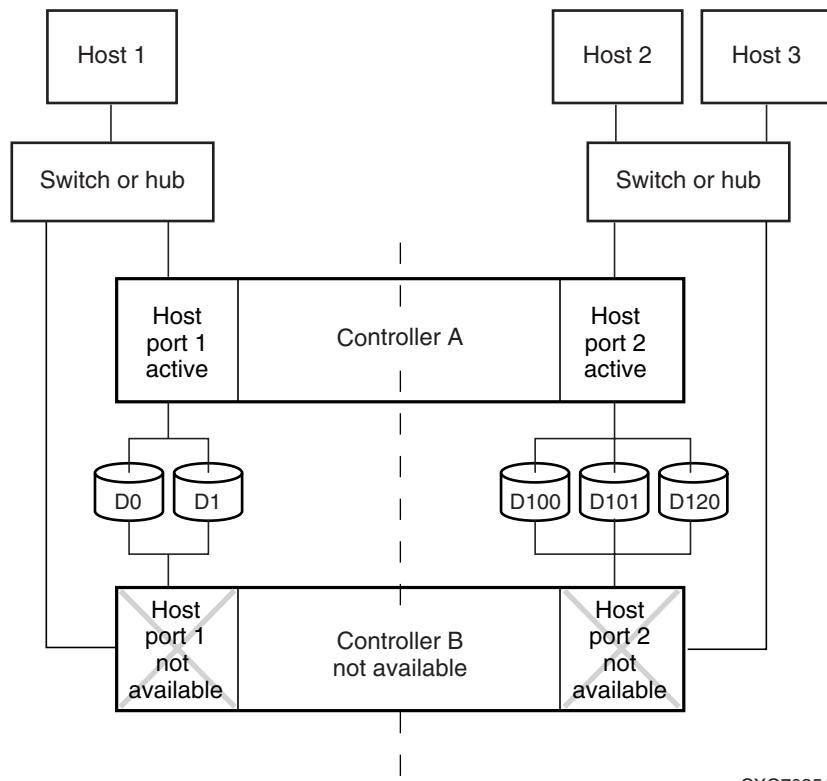


Figure 9: Transparent failover—after failover from controller B to controller A

## Multiple-Bus Failover Mode

Multiple-bus failover mode has the following characteristics:

- Host controls the failover process by moving the units from one controller to another
- All units (0 through 199) are visible at all host ports
- Each host has two or more paths to the units

All hosts must have operating system support or additional software that supports multiple-bus failover mode. With this software, the host sees the same units visible through two (or more) paths. When one path fails, the host can issue commands to move the units from one path to another. A typical multiple-bus failover configuration is shown in [Figure 10](#).

In multiple-bus failover mode, you can specify which units are normally serviced by a specific controller of a controller pair. Units can be preferred to one controller or the other by the PREFERRED\_PATH switch of the ADD UNIT (or SET unit) command. For example, use the following command to prefer unit D101 to “this controller”:

```
SET D101 PREFERRED_PATH=THIS_CONTROLLER
```

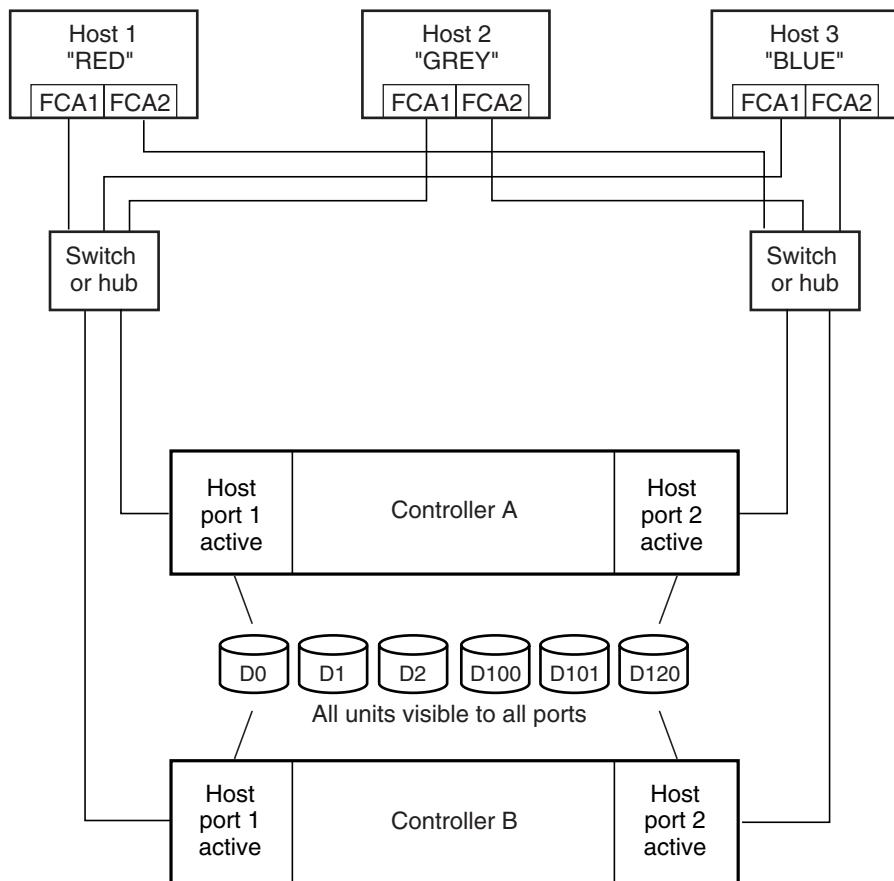
---

**Note:** This is an initial preference, which can be overridden by the hosts.

---

Keep the following points in mind when configuring controllers for multiple-bus failover:

- Multiple-bus failover can compensate for a failure in any of the following:
  - Controller
  - Switch or hub
  - Fibre Channel link
  - Host Fibre Channel adapter
- A host can redistribute the I/O load between the controllers.
- All hosts must have operating system software that supports multiple-bus failover mode.



NOTE: FCA = Fibre Channel Adapter

CXO7094B

**Figure 10: Typical multiple-bus configuration**

## Selecting a Cache Mode

The cache module supports read, read-ahead, write-through, and write-back caching techniques. The cache technique is selected separately for each unit. For example, you can enable only read and write-through caching for some units while enabling only write-back caching for other units.

### Read Caching

When the controller receives a read request from the host, it reads the data from the disk drives, delivers it to the host, and stores the data in its cache module. Subsequent reads for the same data will take the data from cache rather than accessing the data from the disks. This process is called read caching.

Read caching can improve response time to many of the host's read requests. By default, read caching is enabled for all units.

### Read-Ahead Caching

During read-ahead caching, the controller anticipates subsequent read requests and begins to prefetch the next blocks of data from the disks as it sends the requested read data to the host. This is a parallel action. The controller notifies the host of the read completion, and subsequent sequential read requests are satisfied from the cache memory. By default, read-ahead caching is enabled for all units.

### Write-Back Caching

Write-back caching improves the subsystem's response time to write requests by allowing the controller to declare the write operation complete as soon as the data reaches cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.

By default, write-back caching is enabled for all units, but only if there is a backup power source for the cache modules (either batteries or an uninterruptable power supply).

### Write-Through Caching

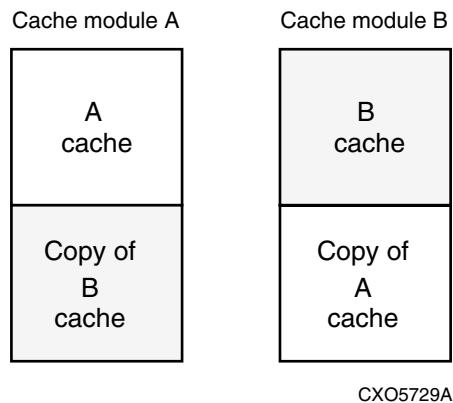
Write-through caching is enabled when write-back caching is disabled. When the controller receives a write request from the host, it places the data in its cache module, writes the data to the disk drives, then notifies the host when the write

operation is complete. This process is called write-through caching because the data actually passes through—and is stored in—the cache memory on its way to the disk drives.

## Enabling Mirrored Caching

In mirrored caching, half of each controller's cache mirrors the companion controller's cache, as shown in [Figure 11](#).

The total memory available for cached data is reduced by half, but the level of protection is greater.



CXO5729A

**Figure 11: Mirrored caching**

Before enabling mirrored caching, make sure the following conditions are met:

- Both controllers support the same size cache.
- Diagnostics indicate that both caches are good.
- No unit errors are outstanding, for example, lost data or data that cannot be written to devices.
- Both controllers are started and configured in failover mode.

## What is the Command Console LUN?

StorageWorks Command Console (SWCC) software communicates with the HSG80 controllers through an existing storage unit, or logical unit number (LUN). The dedicated LUN that SWCC uses is called the Command Console LUN (CCL). The CCL serves as the communication device for the HS-Series Agent and identifies itself to the host by a unique identification string. By default, a CCL device is enabled within the HSG80 controller on host port 1. The HSG80 both SCSI-2 or SCSI-3 uses with your platform.

The CCL does the following:

- Allows the RAID Array to be recognized by the host as soon as it is attached to the SCSI bus and configured into the operating system.
- Serves as a communications device for the HS-Series Agent. The CCL identifies itself to the host by a unique identification string.

In dual-redundant controller configurations, the commands described in the following sections alter the setting of the CCL on both controllers. The CCL is enabled only on host port 1. At least one storage device of any type must be configured on host port 2 before installing the Agent on a host connected to host port 2.

Select a storage set that you plan to configure and that is not likely to change. This storage set can be used by the Agent to communicate with the RAID Array. Deleting this storage set (LUN) later breaks the connection between the Agent and the RAID Array.

## Determining the Address of the CCL

CCL is enabled by default. Its address can be determined by using the following CLI command:

```
HSG80 > SHOW THIS_CONTROLLER
```

## Enabling/Disabling the CCL in SCSI-2 Mode

### Enabling the CCL

To enable the CCL in SCSI-2 mode, use the following CLI command:

```
HSG80 > SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

## Disabling the CCL

To disable the CCL in SCSI-2 mode, enter the following CLI command:

```
HSG80 > SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN
```

To see the state of the CCL, use the SHOW THIS\_CONTROLLER/ OTHER\_CONTROLLER command. Because the CCL is not an actual LUN, the SHOW UNITS command will not display the CCL location.

## Enabling/Disabling CCL in SCSI-3 Mode

The CCL is enabled all the time in SCSI-3 mode. There is no option to enable/disable.

## Determining Connections

The term “connection” applies to every path between a Fibre Channel adapter in a host computer and an active host port on a controller.

---

**Note:** In ACS Version 8.8-1, the maximum number of supported connections is 96.

---

## Naming Connections

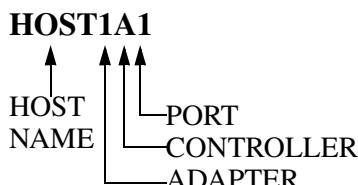
It is highly recommended that you assign names to connections that have meaning in the context of your particular configuration.

---

**Note:** Some non-alpha / numeric characters may not work with some hosts.

---

One system that works well is to name each connection after its host, its adapter, its controller, and its controller host port, as follows:



Examples:

A connection from the first adapter in the host named RED that goes to port 1 of controller A would be called RED1A1.

A connection from the third adapter in host GREEN that goes to port 2 of controller B would be called GREEN3B2.

---

**Note:** Connection names can have a maximum of 9 characters.

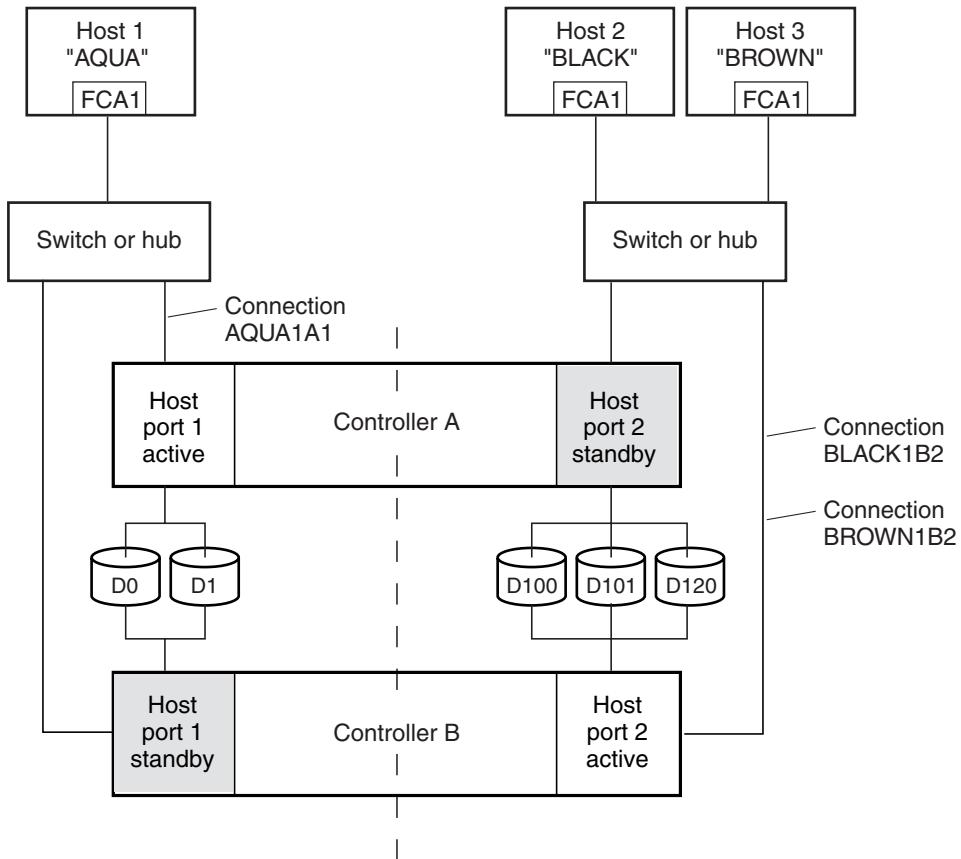
---

## Numbers of Connections

The number of connections resulting from cabling one adapter into a switch or hub depends on failover mode and how many links the configuration has:

- If a controller pair is in transparent failover mode and the port 1 link is separate from the port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another), each adapter will have one connection, as shown in [Figure 12](#).
- If a controller pair is in transparent failover mode and port 1 and port 2 are on the same link (that is, all ports are on the same loop or fabric), each adapter will have two connections, as shown in [Figure 13](#).

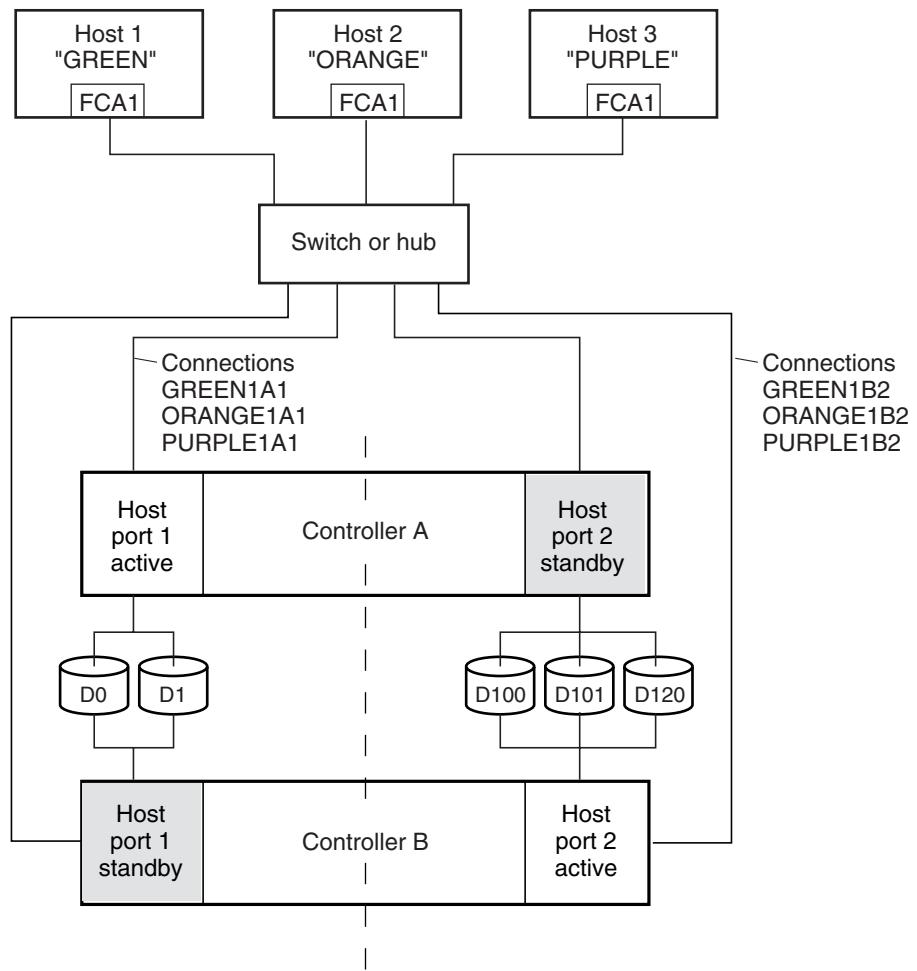
If a controller pair is in multiple-bus failover mode, each adapter has two connections, as shown in [Figure 14](#).



NOTE: FCA = Fibre Channel Adapter

CXO7081B

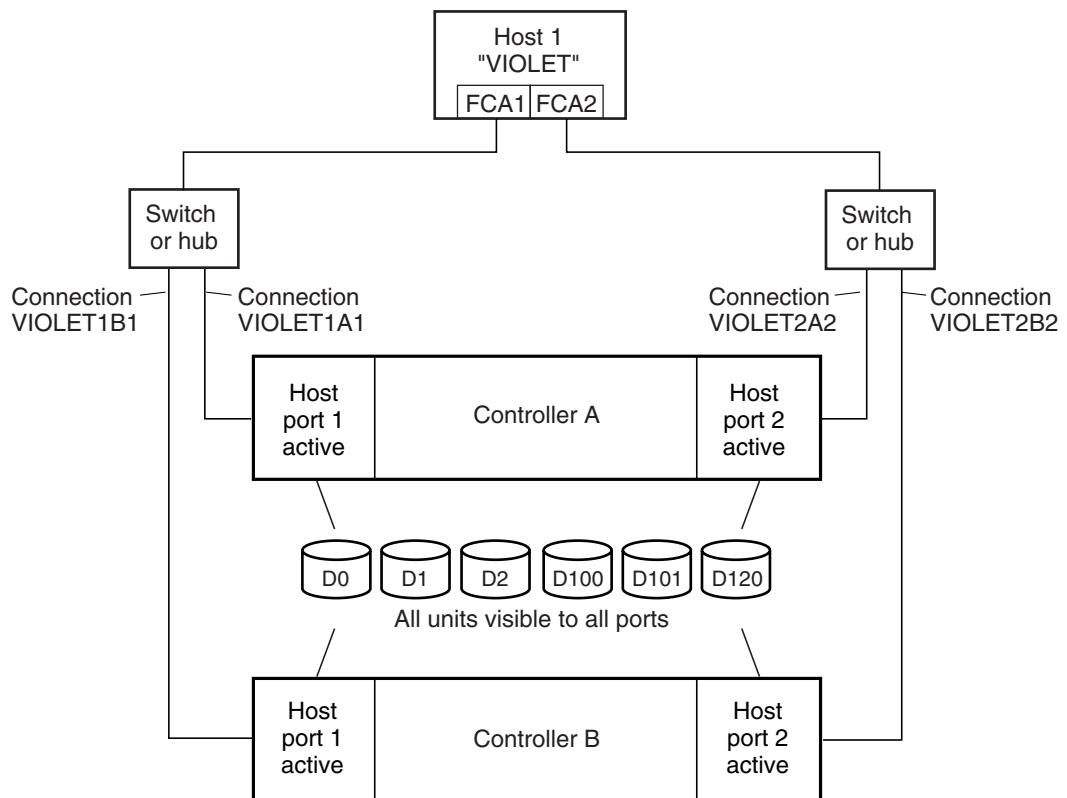
**Figure 12: Connections in separate-link, transparent failover mode configurations**



NOTE: FCA = Fibre Channel Adapter

CXO7079B

**Figure 13: Connections in single-link, transparent failover mode configurations**



NOTE: FCA = Fibre Channel Adapter

CXO7080B

**Figure 14: Connections in multiple-bus failover mode**

## Assigning Unit Numbers

The controller keeps track of the unit with the unit number. The unit number can be from 0–199 prefixed by a D, which stands for disk drive. A unit can be presented as different LUNs to different connections. The interaction of a unit and a connection is determined by several factors:

- Failover mode of the controller pair
- The ENABLE\_ACCESS\_PATH and PREFERRED\_PATH switches in the ADD UNIT (or SET unit) commands
- The UNIT\_OFFSET switch in the ADD CONNECTIONS (or SET connections) commands
- The controller port to which the connection is attached
- The SCSI\_VERSION switch of the SET THIS\_CONTROLLER/OTHER\_CONTROLLER command

The considerations for assigning unit numbers are discussed in the following sections.

## Matching Units to Host Connections in Transparent Failover Mode

In transparent failover mode, the ADD UNIT command creates a unit for host connection to access and assigns it to either port 1 or port 2 of both controllers.

Unit numbers are assigned to ports as follows:

- 0–99 are assigned to host port 1 of both controllers.
- 100–199 are assigned to host port 2 of both controllers.

For example, unit D2 is on port 1, and unit D102 is available through port 2.

The LUN number that a host connection assigns to a unit is a function of the UNIT\_OFFSET switch of the ADD (or SET) CONNECTIONS command. The relationship of offset, LUN number, and unit number is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

Where...

- LUN number is relative to the host (what the host sees the unit as)
- Unit number is relative to the controller (what the controller sees the unit as)

If no value is specified for offset, then connections on port 1 have a default offset of 0 and connections on port 2 have a default offset of 100.

For example, if all host connections use the default offset values, unit D2 will be presented to a port 1 host connection as LUN 2 (unit number of 2 minus offset of 0). Unit D102 will be presented to a port 2 host connection as LUN 2 (unit number of D102 minus offset of 100).

**Figure 15** shows how units are presented as different LUNs, depending on the offset of the host. In this illustration, host connection 1 and host connection 2 would need to be on host port 1; host connection 3 would need to be on host port 2.

Controller units	Host connection 1 Offset: 0	Host connection 2 Offset: 20	Host connection 3 Offset: 100
D0	LUN 0		
D1	LUN 1		
D2	LUN 2		
D3	LUN 3		
D20	LUN 20	LUN 0	
D21	LUN 21	LUN 1	
D100			LUN 0
D101			LUN 1
D102			LUN 2
D130			LUN 30
D131			LUN 31

CXO6455B

**Figure 15: LUN presentation to hosts, as determined by offset**

Offsets other than the default values can be specified. For example, unit D17 would be visible to a host connection on port 1 that had an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible to a host connection with a unit offset of 18 or greater, because that offset is not within the unit's range (unit number of 17 minus offset of 18 is a negative number).

Similarly, unit D127 would be visible to a host connection on port 2 that had an offset of 120 as LUN 7 (unit number of 127 minus offset of 120). The unit would not be visible to a host connection with a unit offset of 128 or greater, because that offset is not within the unit's range (unit number of 127 minus offset of 128 is a negative number).

An additional factor to consider when assigning unit numbers and offsets is SCSI version. If the SCSL\_VERSION switch of the SET THIS\_CONTROLLER/OTHER\_CONTROLLER command is set to SCSI-3, the CCL is presented as LUN 0 to every connection, superseding any unit assignments. The interaction between SCSI version and unit numbers is explained further in the next section.

In addition, the access path to the host connection must be enabled for the connection to access the unit. See “[Restricting Host Access in Transparent Failover Mode](#)” on page 49.

## Matching Units to Host Connections in Multiple-Bus Failover Mode

In multiple-bus failover mode, the ADD UNIT command creates a unit for host connections to access. All unit numbers (0 through 199) are potentially visible on all four controller ports, but are accessible only to those host connections for which access path is enabled and which have offsets in the unit's range.

The LUN number a host connection assigns to a unit is a function of the UNIT\_OFFSET switch of the ADD (or SET) CONNECTIONS command. The default offset is 0. The relationship of offset, LUN number, and unit number is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

Where...

- LUN number is relative to the host (number the host sees the unit as)
- Unit number is relative to the controller (number the controller sees the unit as)

For example, unit D7 would be visible to a host connection with an offset of 0 as LUN 7 (unit number of 7 minus offset of 0). Unit D17 would be visible to a host connection with an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible at all to a host connection with a unit offset of 18 or greater, because that offset is not within the units range (unit number of 17 minus offset of 18 is a negative number).

In addition, the access path to the host connection must be enabled for the connection to access the unit. This is done through the ENABLE\_ACCESS\_PATH switch of the ADD UNIT (or SET unit) command.

The PREFERRED\_PATH switch of the ADD UNIT (or SET unit) command determines which controller of a dual-redundant pair initially accesses the unit. Initially, PREFERRED\_PATH determines which controller presents the unit as Ready. The other controller presents the unit as Not Ready. Hosts can issue a SCSI Start Unit command to move the unit from one controller to the other.

## Assigning Unit Numbers Depending on SCSI\_VERSION

The SCSI\_VERSION switch of the SET THIS\_CONTROLLER/OTHER\_CONTROLLER command determines how the CCL is presented. There are two choices: SCSI-2 and SCSI-3. The choice for SCSI\_VERSION affects how certain unit numbers and certain host connection offsets interact.

### Assigning Host Connection Offsets and Unit Numbers in SCSI-3 Mode

If SCSI\_VERSION is set to SCSI-3, the CCL is presented as LUN 0 to all connections. The CCL supersedes any other unit assignment. Therefore, in SCSI-3 mode, a unit that would normally be presented to a connection as LUN 0 is not visible to that connection at all.

The following methods are recommended for assigning host connection offsets and unit numbers in SCSI-3 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should not be assigned at connection offsets (to avoid being masked by the CCL at LUN 0).

For example, if a host connection has an offset of 20 and SCSI-3 mode is selected, the connection will see LUNs as follows:

LUN 0 - CCL  
LUN 1 - unit 21  
LUN 2 - unit 22, etc.

In this example, if a unit 20 is defined, it will be superseded by the CCL and invisible to the connection.

### Assigning Host Connection Offsets and Unit Numbers in SCSI-2 Mode

Some operating systems expect or require a disk unit to be at LUN 0. In this case, it is necessary to specify SCSI-2 mode.

If SCSI\_VERSION is set to SCSI-2 mode, the CCL floats, moving to the first available LUN location, depending on the configuration.

StorageWorks recommends to use the following conventions when assigning host connection offsets and unit numbers in SCSI-2 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should be assigned at connection offsets (so that every host connection has a unit presented at LUN 0).

**Table 3** summarizes the recommendations for unit assignments based on the `SCSI_VERSION` switch.

**Table 3: Unit Assignments and SCSI\_VERSION**

<code>SCSI_VERSION</code>	<code>Offset</code>	<code>Unit Assignment</code>	<code>What the connection sees LUN 0 as</code>
SCSI-2	Divisible by 10	At offsets	Unit whose number matches offset
SCSI-3	Divisible by 10	Not at offsets	CCL

---

## What is Selective Storage Presentation?

Selective Storage presentation is a feature of the HSG80 controller that enables the user to control the allocation of storage space and shared access to storage across multiple hosts. This is also known as *Restricting Host Access*.

In a subsystem that is attached to more than one host or if the hosts have more than one adapter, it is possible to reserve certain units for the exclusive use of certain host connections.

For a controller pair, the method used to restrict host access depends on whether the controllers are in transparent or multiple-bus failover mode. For a single controller, the methods are the same as for a controller pair in transparent failover.

The default condition is `ENABLE_ACCESS_PATH=ALL`. This specifies that access paths to ALL hosts are enabled. HP recommends that you restrict host access and that the access path be carefully specified to avoid providing undesired host connections access to the unit.

Now with V88, you may change the default subsystem behavior so that units are always created without enabling connections. This provides more control in granting appropriate access to specific connections.

The command syntax to disable is:

```
HSG> SET this/other Default_Access=Disable
```

The command syntax to enable is:

```
HSG> SET this/other Default_Access=Enable {default after  
upgrade}
```

When the command is invoked from one controller, the Default\_Access from the other controller will be similarly modified. The setting is symmetrical and persistent across restarts, FRUTIL, etc.

## Restricting Host Access in Transparent Failover Mode

Three methods can be used to restrict host access to storage units in transparent failover mode:

- Using separate Fibre Channel links (either loop or fabric)
- Enabling the access path of selected host connections on a shared loop or fabric
- Setting offsets

---

**Note:** These techniques also work for a single controller.

---

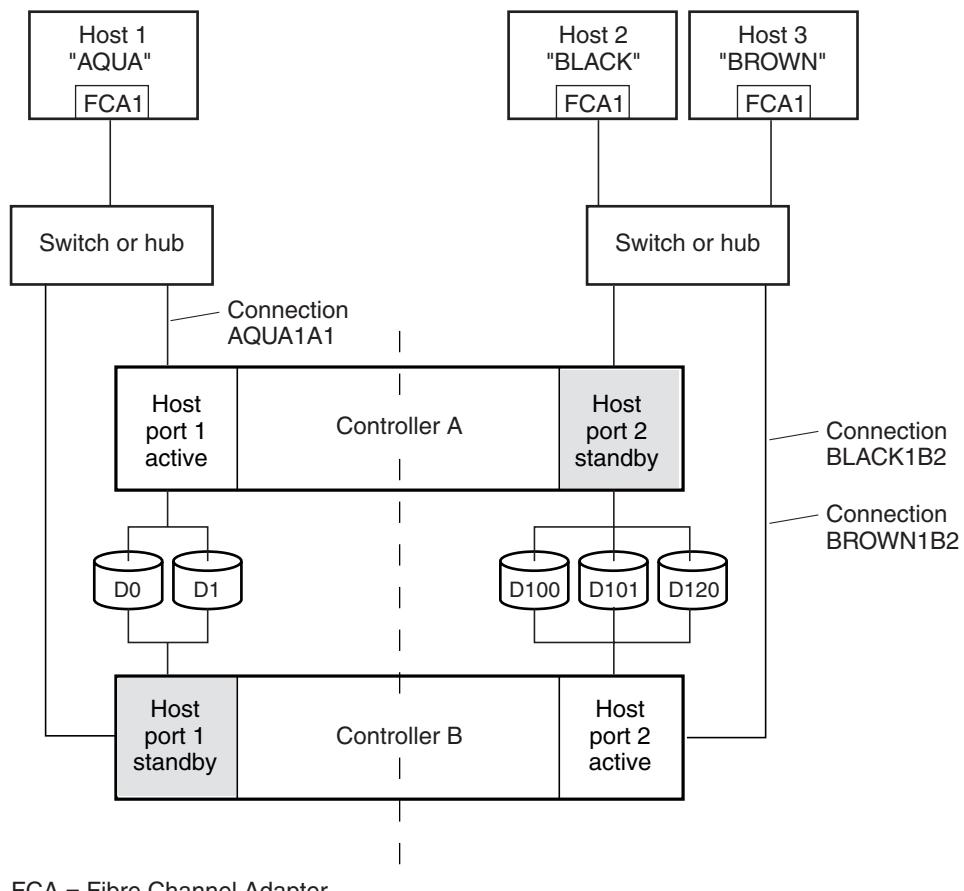
### Restricting Host Access by Separate Links

In transparent failover mode, host port 1 of controller A and host port 1 of controller B share a common Fibre Channel link. Host port 2 of controller A and host port 2 of controller B also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way to limit host access is to have one host or set of hosts on the port 1 link, and another host or set of hosts on the port 2 link. Each host can then see only units assigned to its respective controller port. This separation of host buses is shown in [Figure 16](#). This method applies only if the host 1 link and host 2 link are separate links.

---

**Note:** It is highly recommended that you provide access to only specific connections. This way, if new connections are added, they will not have automatic access to all units. See “[Restricting Host Access by Disabling Access Paths](#)” on page 50.

---



FCA = Fibre Channel Adapter

CXO7081B

Figure 16: Limiting host access in transparent failover mode

## Restricting Host Access by Disabling Access Paths

If more than one host is on a link (that is, attached to the same port), host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the `ENABLE_ACCESS_PATH` and `DISABLE_ACCESS_PATH` switches of the `ADD UNIT` (or `SET unit`) commands. The access path is a unit switch, meaning it must be specified for each unit. Default access enables the unit to be accessible to all hosts.

For example:

In [Figure 17](#), restricting the access of unit D101 to host 3, the host named BROWN can be done by enabling only the connection to host 3. Enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=BROWN1B2
```

If the storage subsystem has more than one host connection, carefully specify the access path to avoid providing undesired host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify DISABLE\_ACCESS\_PATH=ALL for the unit, then specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

---

**Note:** The procedure of restricting access by enabling all access paths then disabling selected paths is not recommended because of the potential data/security breach that occurs when a new host connection is added.

---

## Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access.

For example:

In Figure 16, assume both host connections on port 2 (connections BLACK1B2 and BROWN1B2) initially have the default port 2 offset of 100. Setting the offset of connection BROWN1B2 to 120 will present unit D120 to host BROWN as LUN 0.

```
SET BROWN1B2 UNIT_OFFSET=120
```

Host BROWN cannot see units lower than its offset, so it cannot access units D100 and D101. However, host BLACK can still access D120 as LUN 20 if the operating system permits. To restrict access of D120 to only host BROWN, enable only host BROWN's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL  
SET D120 ENABLE_ACCESS_PATH=BROWN1B2
```

---

**Note:** StorageWorks recommends that you provide access to only specific connections, even if there is just one connection on the link. This way, if new connections are added, they will not have automatic access to all units.

---

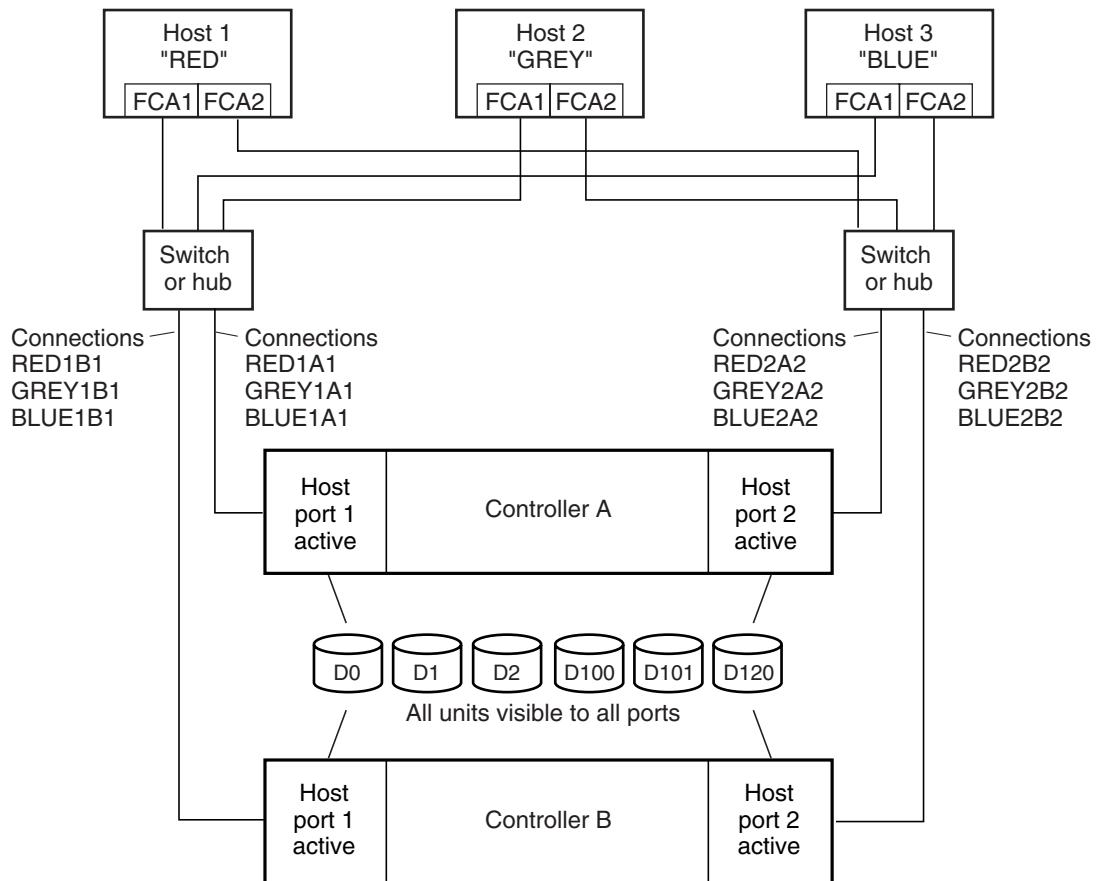
## Restricting Host Access in Multiple-Bus Failover Mode

In multiple-bus mode, the units assigned to any port are visible to all ports. There are two ways to limit host access in multiple-bus failover mode:

- Enabling the access path of selected host connections
- Setting offsets

### Enabling the Access Path of Selected Host Connections

Host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the ENABLE\_ACCESS\_PATH and DISABLE\_ACCESS\_PATH switches of the ADD UNIT (or SET unit) commands. Access path is a unit switch, meaning it must be specified for each unit. Default access means that the unit is accessible to all hosts. It is important to remember that at least two paths between the unit and the host must be enabled in order for multiple-bus failover to work.



NOTE: FCA = Fibre Channel Adapter

CXO7078B

**Figure 17: Limiting host access in multiple-bus failover mode**

For example:

Figure 17 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host BLUE can be done by enabling only the connections to host BLUE. At least two connections must be enabled for

multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. To enable all connections for host BLUE, enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=BLUE1A1,BLUE1B1,BLUE2A2,BLUE2B2
```

To enable only two connections for host BLUE (if it is a restriction of the operating system), select two connections that use different adapters, different switches or hubs, and different controllers:

```
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE2B2)  
or  
SET D101 DISABLE_ACCESS_PATH=ALL  
SET D101 ENABLE_ACCESS_PATH=(BLUE1B1, BLUE2A2)
```

If the storage subsystem has more than one host connection, the access path must be specified carefully to avoid giving undesirable host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify DISABLE\_ACCESS\_PATH=ALL when the unit is added, then use the *SET unit* command to specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

---

**Note:** The procedure of restricting access by enabling all access paths then disabling selected paths is not recommended because of the potential data/security breach that occurs when a new host connection is added.

---

## Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access. However, depending on the operating system, hosts that have lower offsets may be able to access the units in the specified range.

---

**Note:** All host connections to the same host computer must be set to the same offset.

---

For example:

In [Figure 17](#), assume all host connections initially have the default offset of 0. Giving all connections access to host BLUE, an offset of 120 will present unit D120 to host BLUE as LUN 0. Enter the following commands:

```
SET BLUE1A1 UNIT_OFFSET=120
SET BLUE1B1 UNIT_OFFSET=120
SET BLUE2A2 UNIT_OFFSET=120
SET BLUE2B2 UNIT_OFFSET=120
```

Host BLUE cannot see units lower than its offset, so it cannot access any other units. However, the other two hosts can still access D120 as LUN 20 if their operating system permits. To restrict access of D120 to only host BLUE, enable only host BLUE's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(BLUE1A1,BLUE1B1,BLUE2A2,BLUE2B2)
```

---

**Note:** StorageWorks recommends that you always provide access to only specific connections. This way, if new connections are added, they will not have automatic access to all units.

---

---

**Note:** See [“Restricting Host Access by Disabling Access Paths”](#) on page 50.

---

## Worldwide Names (Node IDs and Port IDs)

A worldwide name—also called a node ID—is a unique, 64-bit number assigned to a subsystem prior to shipping. The node ID belongs to the subsystem itself and never changes.

Each subsystem's node ID ends in zero, for example 5000-1FE1-FF0C-EE00. The controller port IDs are derived from the node ID.

In a subsystem with two controllers in transparent failover mode, the controller port IDs are incremented as follows:

- Controller A and controller B, port 1—worldwide name + 1, for example 5000-1FE1-FF0C-EE01
- Controller A and controller B, port 2—worldwide name + 2, for example 5000-1FE1-FF0C-EE02

In multiple-bus failover mode, each of the host ports has its own port ID:

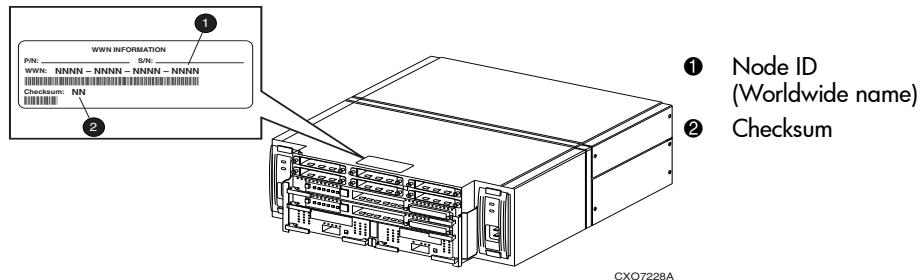
- Controller B, port 1—worldwide name + 1, for example 5000-1FE1-FF0C-EE01
- Controller B, port 2—worldwide name + 2, for example 5000-1FE1-FF0C-EE02
- Controller A, port 1—worldwide name + 3, for example 5000-1FE1-FF0C-EE03
- Controller A, port 2—worldwide name + 4, for example 5000-1FE1-FF0C-EE04

Use the CLI command, `SHOW THIS_CONTROLLER/OTHER_CONTROLLER` to display the subsystem's worldwide name.

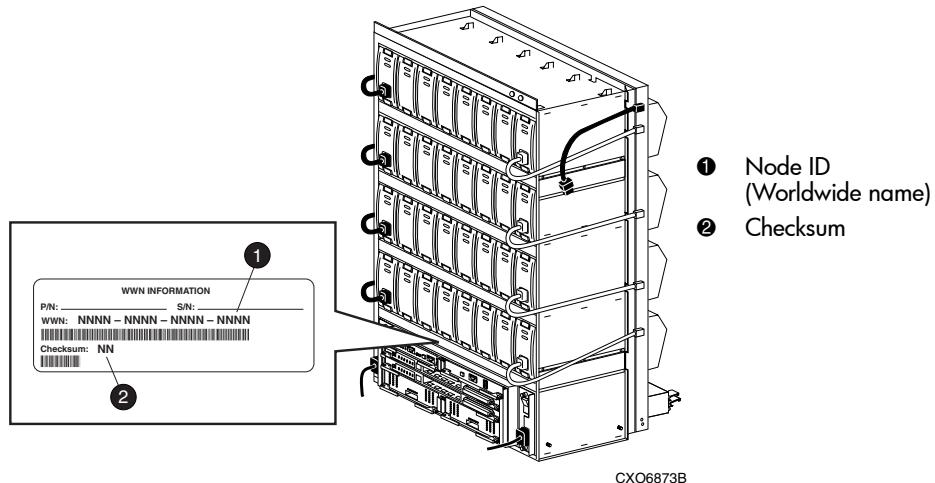
## Restoring Worldwide Names (Node IDs)

If a situation occurs that requires you to restore the worldwide name, you can restore it using the worldwide name and checksum printed on the sticker on the frame into which the controller is inserted.

[Figure 18](#) shows the placement of the worldwide name label for the Model 2200 enclosure, and [Figure 19](#) for the BA370 enclosure.



**Figure 18: Placement of the worldwide name label on the Model 2200 enclosure**



**Figure 19: Placement of the worldwide name label on the BA370 enclosure**



**Caution:** Each subsystem has its own unique worldwide name (node ID). If you attempt to set the subsystem worldwide name to a name other than the one that came with the subsystem, the data on the subsystem will not be accessible. Never set two subsystems to the same worldwide name, or data corruption will occur.

## Unit Worldwide Names (LUN IDs)

In addition, each unit has its own worldwide name, or LUN ID. This is a unique, 128-bit value that the controller assigns at the time of unit initialization. It cannot be altered by the user but does change when the unit is reinitialized. Use the SHOW command to list the LUN ID.



# Planning Storage Configurations

2

This chapter provides information to help you plan the storage configuration of your subsystem. Storage containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. Use the guidelines found in this section to plan the various types of storage containers needed.

This chapter also focuses on the required design and implementation aspects, such as addressing conventions, configuration rules, creating storage profiles, and creating storage maps.

The following storage configuration information can be found in this chapter:

- [Where to Start](#), page 60
- [Determining Storage Requirements](#), page 61
- [Configuration Rules for the Controller](#), page 61
- [Addressing Conventions for Device PTL](#), page 62
- [Choosing a Container Type](#), page 72
- [Creating a Storageset Profile](#), page 75
- [Planning Considerations for Storageset](#), page 76
- [Changing Characteristics Through Switches](#), page 87
- [Specifying Storageset and Partition Switches](#), page 88
- [Specifying Initialization Switches](#), page 89
- [Specifying Unit Switches](#), page 93
- [Creating Storage Maps](#), page 93

Refer to Chapter 3 for instructions on how to prepare your host computer to accommodate the HSG80 controller storage subsystem.

## Where to Start

The following procedure outlines the steps to follow when planning your storage configuration.

See Appendix A to locate the blank templates for keeping track of the containers being configured.

1. Determine your storage requirements. Use the questions in “[Determining Storage Requirements](#)” on page 61, to help you.
2. Review configuration rules. See “[Configuration Rules for the Controller](#)” on page 61.
3. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See “[Addressing Conventions for Device PTL](#)” on page 62.
4. Choose the type of storage containers you need to use in your subsystem. See “[Choosing a Container Type](#)” on page 72, for a comparison and description of each type of storageset.
5. Create a storageset profile (described in “[Creating a Storageset Profile](#)” on page 75). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
  - [Planning Considerations for Storageset](#), page 76
  - [Mirrorset Planning Considerations](#), page 80
  - [RAIDset Planning Considerations](#), page 81
  - [Partition Planning Considerations](#), page 85
  - [Striped Mirrorset Planning Considerations](#), page 83
6. Decide which switches you need for your subsystem. General information on switches is detailed in “[Specifying Storageset and Partition Switches](#)” on page 88.
  - Determine the unit switches you want for your units (“[Specifying Unit Switches](#)” on page 93).
  - Determine the initialization switches you want for your planned storage containers (“[Specifying Initialization Switches](#)” on page 89).
7. Create a storage map (“[Creating Storage Maps](#)” on page 93).
8. Configure the storage you have now planned using one of the following methods:

- Use SWCC. See the SWCC documentation for details.
- Use the Command Line Interpreter (CLI) commands. This method allows you flexibility in defining and naming your storage containers. See the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

## Determining Storage Requirements

It is important to determine your storage requirements. Here are a few of the questions you should ask yourself regarding the subsystem usage:

- What applications or user groups will access the subsystem? How much capacity do they need?
- What are the I/O requirements? If an application is data transfer-intensive, what is the required transfer rate? If it is I/O request-intensive, what is the required response time? What is the read/write ratio for a typical request?
- Are most I/O requests directed to a small percentage of the disk drives? Do you want to keep it that way or balance the I/O load?
- Do you store mission-critical data? Is availability the highest priority or would standard backup procedures suffice?

## Configuration Rules for the Controller

The following list defines maximum configuration rules for the controller:

- 128 visible LUNs/200 assignable unit numbers
  - In SCSI-2 mode, if the CCL is enabled, the result is 127 visible LUNs and one CCL.
  - In SCSI-3 mode, if the CCL is enabled, the result is 126 visible LUNs and two CCLs.
- 1.024 TB storage set size
- 96 host connections
- 84 physical devices
- 20 RAID 3/5 storage sets
- 30 RAID 3/5 and RAID 1 storage sets (see note)
- 45 RAID 3/5, RAID 1, and RAID 0 storage sets (see note)

---

**Note:** For the previous two storageset configurations, this is a combined maximum, limited to no more than 20 RAID 3/5 storagesets in the individual combination.

---

- 8 partitions of a storageset or individual disk
- 6 physical devices per RAID 1 storageset (mirrorset)
- 14 physical devices per RAID 3/5 storageset (RAIDset)
- 24 physical devices per RAID 0 storageset (stripeset)
- 48 physical devices per RAID 0+1 storageset (striped mirrorset)
- 146 GB disks—the largest supported by the HSG60 and HSG80 storage array.

---

**Tip:** If you are redeploying disks that have been operating under a prior version of ACS into a newly established container, as a best practice, *always* initialize the devices and the new container before proceeding with subsystem activities to avoid operational and performance issues.

---

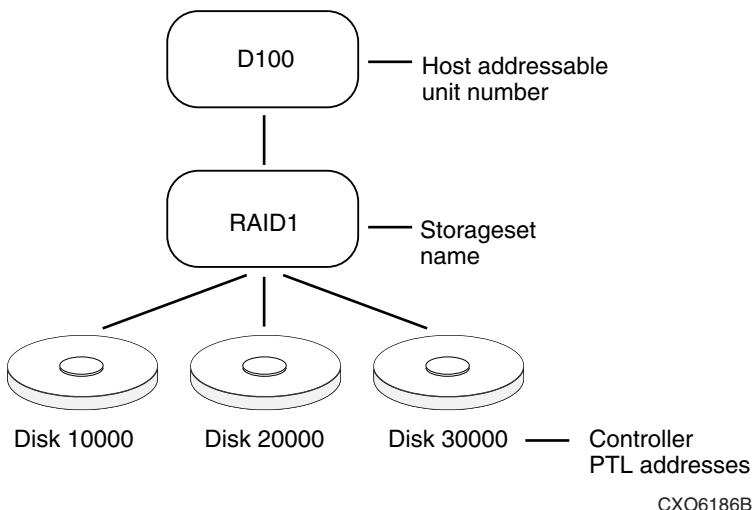
---

**Note:** HSG60 and HSG80 array controllers do not support greater than 146 GB. This restriction is enforced by the firmware.

---

## Addressing Conventions for Device PTL

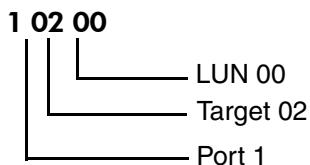
The HSG80 controller has six SCSI device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (The *StorageWorks Command Console* calls the device ports “channels.”) The standard BA370 enclosure provides a maximum of four SCSI target identifications (ID) for each device port. If more target IDs are needed, expansion enclosures can be added to the subsystem. For an example of how units are mapped to physical disk drives, see [Figure 20](#).



**Figure 20: Mapping a unit to physical disk drives**

The HSG80 controller identifies devices based on a Port-Target-LUN (PTL) numbering scheme, shown in [Figure 21](#). The physical location of a device in its enclosure determines its PTL.

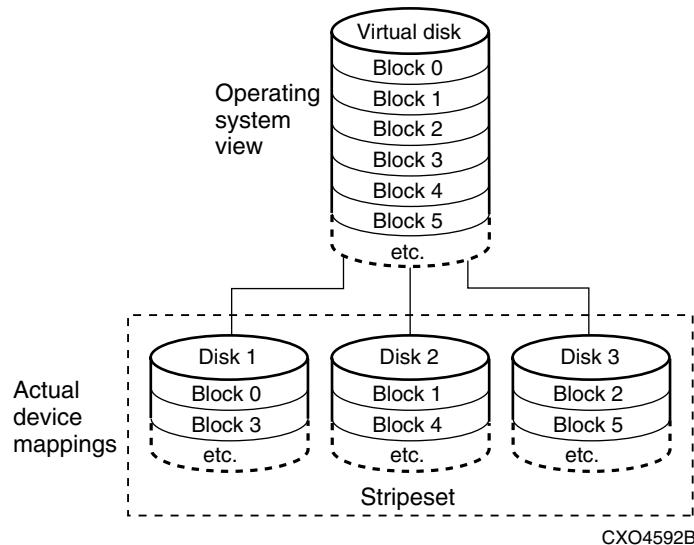
- P—Designates the controller's SCSI device port number (1 through 6).
- T—Designates the target ID number of the device. Valid target ID numbers for a single-controller configuration and dual-redundant controller configuration are 0–3 and 8–15, respectively. (This applies to the BA370 cabinet only.)
- L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.



**Figure 21: PTL naming convention**

The controller can either operate with a BA370 enclosure or with a Model 2200 controller enclosure combined with Model 4214R, Model 4254, Model 4310R, Model 4350R, Model 4314R, or Model 4354R disk enclosures.

The controller operates with BA370 enclosures that are assigned ID numbers 0, 2, and 3. These ID numbers are set through the PVA module. Enclosure ID number 1, which assigns devices to targets 4 through 7, is not supported. [Figure 22](#) shows how data is laid out on disks in an extended configuration.



**Figure 22: How data is laid out on disks in BA370 enclosure configuration**

## Examples - Model 2200 Storage Maps, PTL Addressing

The Model 2200 controller enclosure can be combined with the following:

- **Model 4214R disk enclosure** — Ultra2 SCSI with 14 drive bays, single-bus I/O module.
- **Model 4254 disk enclosure** — Ultra2 SCSI with 14 drive bays, dual-bus I/O module.

---

**Note:** The Model 4214R uses the same storage maps as the Model 4314R, and the Model 4254 uses the same storage maps as the Model 4354R disk enclosures.

---

- **Model 4310R disk enclosure** — Ultra3 SCSI with 10 drive bays, single-bus I/O module. [Table 4](#) shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4310R disk enclosures can be used with each Model 2200 controller enclosure.

---

**Note:** The storage map for the Model 4310R reflects the disk enclosure's physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure, and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

---

- **Model 4350R disk enclosure** — Ultra3 SCSI with 10 drive bays, single-bus I/O module. [Table 5](#) shows the addresses for each device in a three-shelf, single-bus configuration. A maximum of three Model 4350R disk enclosures can be used with each Model 2200 controller enclosure.
- **Model 4314R disk enclosure** — Ultra3 SCSI with 14 drive bays, single-bus I/O module. [Table 6](#) shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4314R disk enclosures can be used with each Model 2200 controller enclosure.

---

**Note:** The storage map for the Model 4314R reflects the disk enclosure's physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure, and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

---

- **Model 4354R disk enclosure** — Ultra3 SCSI with 14 drive bays, dual-bus I/O module. [Table 7](#) shows the addresses for each device in a three-shelf, dual-bus configuration. A maximum of three Model 4354R disk enclosures can be used with each Model 2200 controller enclosure.

---

**Note:** Appendix A contains storageset profiles you can copy and use to create your own system profiles. It also contains an enclosure template you can use to help you keep track of the location of devices and storagesets in your shelves.

---

**Table 4: PTL addressing, single-bus configuration, six Model 4320R enclosures**

Model 4310R Disk Enclosure Shelf 6 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200
Model 4310R Disk Enclosure Shelf 5 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200
Model 4310R Disk Enclosure Shelf 4 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

**Table 4: PTL addressing, single-bus configuration, six Model 4320R enclosures (Continued)**

Model 4310R Disk Enclosure Shelf 1 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk11000	Disk11100	Disk11200
Model 4310R Disk Enclosure Shelf 2 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk21000	Disk21100	Disk21200
Model 4310R Disk Enclosure Shelf 3 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

Table 5: PTL addressing, dual-bus configuration, three Model 4350R enclosures

Model 4350R Disk Enclosure Shelf 1 (Single-Bus)										
	SCSCSI Bus ASI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400
Model 4350R Disk Enclosure Shelf 2 (Single-Bus)										
	SCSCSI Bus ASI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400
Model 4350R Disk Enclosure Shelf 3 (Single-Bus)										
	SCSCSI Bus ASI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400

**Table 6: PTL addressing, single-bus configuration, six Model 4314R enclosures**

Model 4314R Disk Enclosure Shelf 6 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500
Model 4314R Disk Enclosure Shelf 5 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51400	Disk51500
Model 4314R Disk Enclosure Shelf 4 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	Disk41400	Disk41500
Model 4314R Disk Enclosure Shelf 1 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

**Table 6: PTL addressing, single-bus configuration, six Model 4314R enclosures (Continued)**

<b>Model 4314R Disk Enclosure Shelf 2 (Single-Bus)</b>														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500
<b>Model 4314R Disk Enclosure Shelf 3 (Single-Bus)</b>														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

**Table 7: PTL addressing, dual-bus configuration, three Model 4354A enclosures**

<b>Model 4354R Disk Enclosure Shelf 1 (Dual-Bus)</b>														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800
<b>Model 4354R Disk Enclosure Shelf 2 (Dual-Bus)</b>														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08

**Table 7: PTL addressing, dual-bus configuration, three Model 4354A enclosures (Continued)**

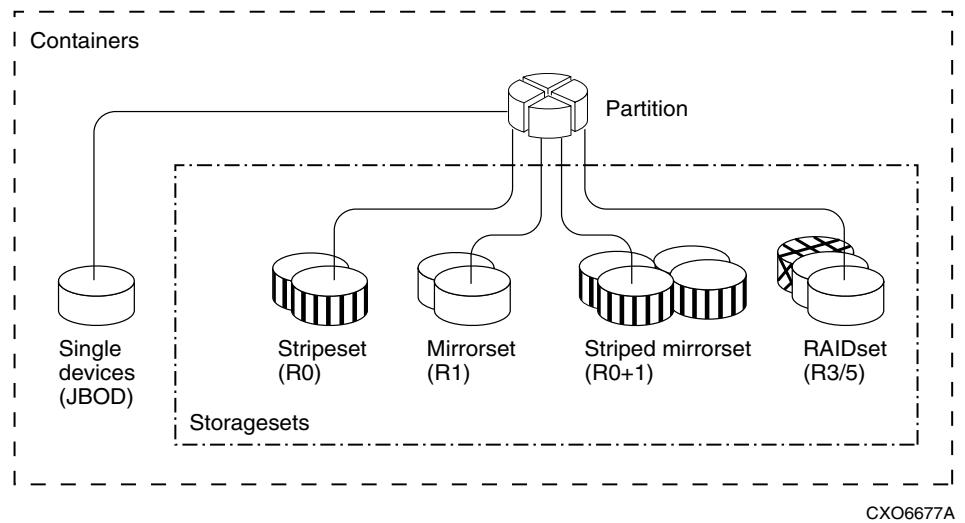
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800
<b>Model 4354R Disk Enclosure Shelf 3 (Dual-Bus)</b>														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

## Choosing a Container Type

Different applications may have different storage requirements. You probably want to configure more than one kind of container within your subsystem.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in [Figure 23](#). The independent disks and the selected storageset may also be partitioned.

The storagesets implement RAID (Redundant Array of Independent Disks) technology. Consequently, they all share one important feature: each storageset, whether it contains two disk drives or ten, looks like one large, virtual disk drive to the host.



**Figure 23: Storage container types**

**Table 8** compares the different kinds of containers to help you determine which ones satisfy your requirements.

**Table 8: Comparison of Container Types**

Container Name	Relative Availability	Request Rate (Read/Write) I/O Per Second	Transfer Rate (Read/Write) MB per second	Applications
Independent disk drives (JBOD)	Equal to number of JBOD disk drives	Comparable to single disk drive	Comparable to single disk drive	—
Stripeset (RAID 0)	Proportionate to number of disk drives; worse than single disk drive	Excellent if used with large chunk size	Excellent if used with small chunk size	High performance for non-critical data
Mirrorset (RAID1)	Excellent	Good/Fair	Good/Fair	System drives; critical files
RAIDset (RAID 3/5)	Excellent	Excellent/good	Read: excellent (if used with small chunk sizes) Write: good (if used with small chunk sizes)	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0+1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

For a comprehensive discussion of RAID, refer to *The RAIDBOOK—A Source Book for Disk Array Technology*.

## Creating a StorageSet Profile

Creating a profile for your storage sets, partitions, and devices can simplify the configuration process. Filling out a storageSet profile helps you choose the storage sets that best suit your needs and to make informed decisions about the switches you can enable for each storageSet or storage device that you configure in your subsystem.

For an example of a storageSet profile, see [Table 9](#). This table contains blank profiles that you can copy and use to record the details for your storage sets. Use the information in this chapter to help you make decisions when creating storageSet profiles.

**Table 9: Example of StorageSet Profile**

Type of StorageSet:							
<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD			
Mirrorset							
<b>StorageSet Name</b> R1. <b>Disk Drives</b> D10300, D20300, D10400, D20400							
<b>Unit Number</b> D101							
Partitions:							
%	%	%	%	%	%	%	%
RAIDset Switches:							
<input checked="" type="checkbox"/> Normal (default)	<input checked="" type="checkbox"/> No (default)	<input checked="" type="checkbox"/> Best performance (default)					
<input type="checkbox"/> Fast	<input type="checkbox"/> Yes, missing:	<input type="checkbox"/> Best fit					
		<input type="checkbox"/> None					
Mirrorset Switches:							
<input type="checkbox"/> Best performance (default)	<input type="checkbox"/> Normal (default)	<input type="checkbox"/> Least busy (default)					
<input type="checkbox"/> Best fit	<input type="checkbox"/> Fast	<input type="checkbox"/> Round robin					
<input type="checkbox"/> None		<input type="checkbox"/> Disk drive:					

## Initialize Switches:

<input checked="" type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input checked="" type="checkbox"/> Destroy (default)
64 blocks	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Retain
128 blocks		
256 blocks		

## Unit Switches:

Read caching	<input checked="" type="checkbox"/> X	<u>ALL</u>
Read-ahead caching		
Write-back caching	<input checked="" type="checkbox"/> X	
Write-through caching		

## Planning Considerations for Stageset

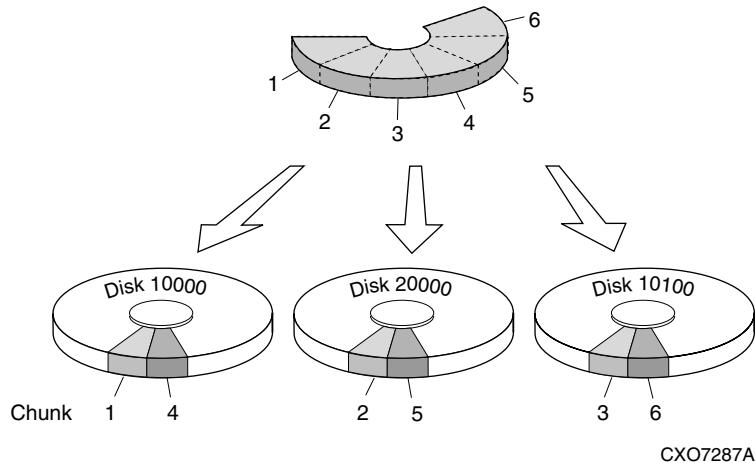
This section contains the guidelines for choosing the stageset type needed for your subsystem:

- [Stripeset Planning Considerations](#), page 76
- [Mirrorset Planning Considerations](#), page 80
- [RAIDset Planning Considerations](#), page 81
- [Striped Mirrorset Planning Considerations](#), page 83
- [Stageset Expansion Considerations](#), page 85
- [Partition Planning Considerations](#), page 85

### Stripeset Planning Considerations

Stripesets (RAID 0) enhance I/O performance by spreading the data across multiple disk drives. Each I/O request is broken into small segments called “chunks.” These chunks are then simultaneously “striped” across the disk drives in the stageset, thereby enabling several disk drives to participate in one I/O request.

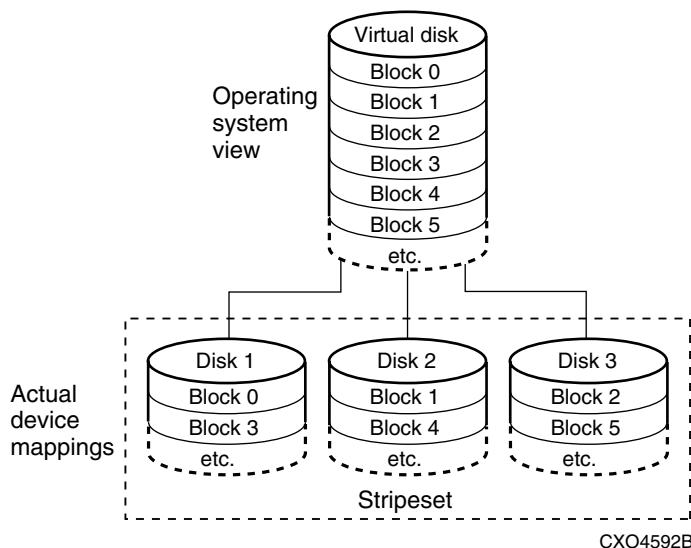
For example, in a three-member stripeset that contains disk drives Disk 10000, Disk 20000, and Disk 10100, the first chunk of an I/O request is written to Disk 10000, the second to Disk 20000, the third to Disk 10100, the fourth to Disk 10000, until all of the data has been written to the drives (Figure 24).



**Figure 24: Three-member RAID 0 stripeset (example 1)**

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rate. You can set the chunk size or use the default setting (see “[Chunk Size](#)” on page 89, for information about setting the chunk size). Figure 25 shows another example of a three-member RAID 0 stripeset.

A major benefit of striping is that it balances the I/O load across all of the disk drives in the storage set. This can increase the subsystem performance by eliminating the hot spots (high localities of reference) that occur when frequently accessed data becomes concentrated on a single disk drive.



**Figure 25: Three-member RAID 0 stripeset (example 2)**

Keep the following points in mind as you plan your stripesets:

- Reporting methods and size limitations prevent certain operating systems from working with large stripesets.
- A storage set should only contain disk drives of the same capacity. The controller limits the effective capacity of each member to the capacity of the smallest member in the storage set (base member size) when the storage set is initialized. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storage set, you waste 5 GB of capacity on each 9 GB member.

If you need high performance and high availability, consider using a RAIDset, striped-mirrorset, or a host-based shadow of a stripeset.

- Striping does not protect against data loss. In fact, because the failure of one member is equivalent to the failure of the entire stripeset, the likelihood of losing data is higher for a stripeset than for a single disk drive.

For example, if the mean time between failures (MTBF) for a single disk is 1 hour, then the MTBF for a stripeset that comprises  $N$  such disks is  $1/N$  hours. As another example, if the MTBF of a single disk is 150,000 hours (about 17 years), a stripeset comprising four of these disks would only have an MTBF of slightly more than 4 years.

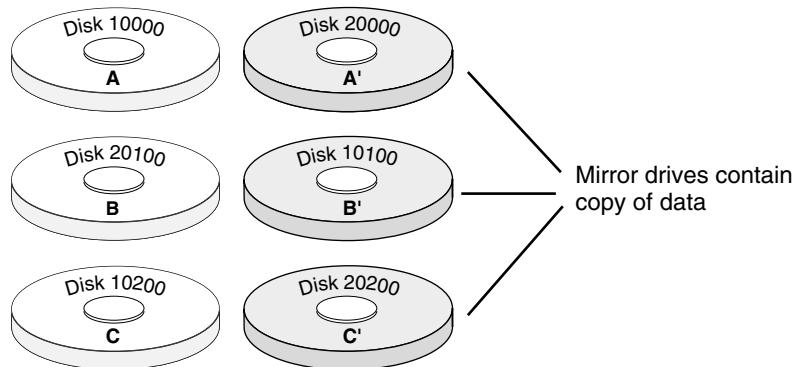
For this reason, you should avoid using a stripeset to store critical data. Stripesets are more suitable for storing data that can be reproduced easily or whose loss does not prevent the system from supporting its critical mission.

- Evenly distribute the members across the device ports to balance the load and provide multiple paths.
- Stripesets may contain between two and 24 members.
- If you plan to use mirror members to replace failing drives, then create the original stripeset as a stripeset of 1-member mirrorsets.
- Stripesets are well-suited for the following applications:
  - Storing program image libraries or run-time libraries for rapid loading.
  - Storing large tables or other structures of read-only data for rapid application access.
  - Collecting data from external sources at very high data transfer rates.
- Stripesets are not well-suited for the following applications:
  - A storage solution for data that cannot be easily reproduced or for data that must be available for system operation.
  - Applications that make requests for small amounts of sequentially located data.
  - Applications that make synchronous random requests for small amounts of data.

Spread the member drives as evenly as possible across the six I/O device ports.

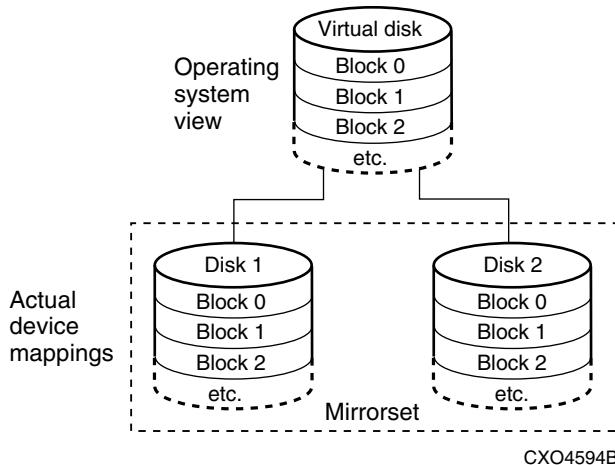
## Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in [Figure 26](#). For each primary disk drive, there is at least one mirror disk drive. Thus, if a primary disk drive fails, its mirror drive immediately provides an exact copy of the data. [Figure 27](#) shows a second example of a Mirrorset.



CXO7288A

**Figure 26: Mirrorsets maintain two copies of the same data (example 1)**



CXO4594B

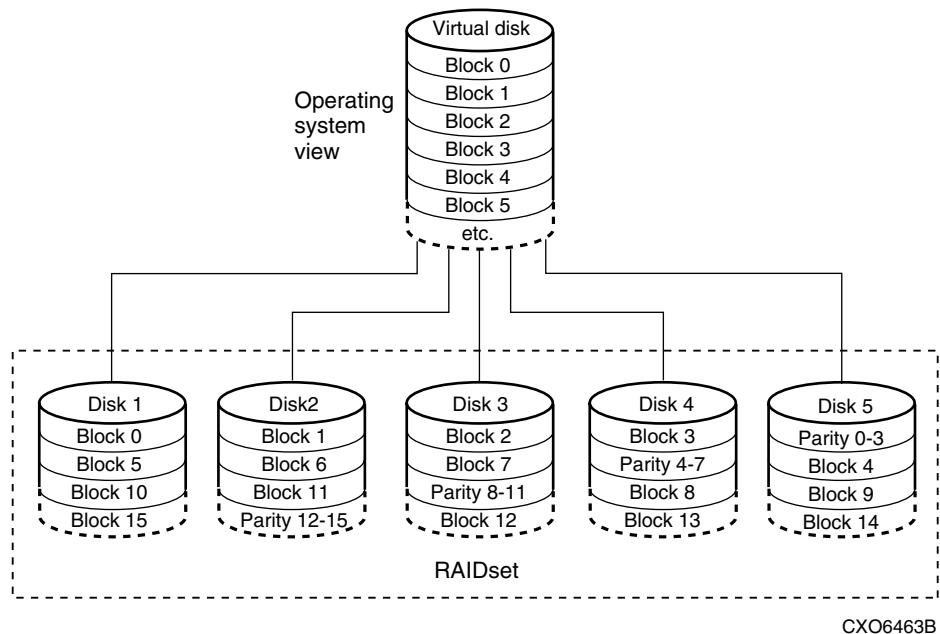
**Figure 27: Mirrorset (example 2)**

## Keep these points in mind when planning mirrorsets

- Data availability with a mirrorset is excellent but comes with a higher cost—you need twice as many disk drives to satisfy a given capacity requirement. If availability is your top priority, consider using dual-redundant controllers and redundant power supplies.
- You can configure up to a maximum of 20 RAID 3/5 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members. Refer to “[Configuration Rules for the Controller](#)” on page 61, for detailed information on maximum numbers. 30 RAID 3/5 and RAID 1 mirrorsets are permitted, however, there is limit of no more than 20 RAID 3/5 mirrorsets in such a configuration.
- Both write-back cache modules must be the same size.
- A mirrorset should only contain disk drives of the same capacity.
- Spread mirrorset members across different device ports (drive bays).
- Mirrorsets are well-suited for the following:
  - Any data for which reliability requirements are extremely high
  - Data to which high-performance access is required
  - Applications for which cost is a secondary issue
- Mirrorsets are not well-suited for the following applications:
  - Write-intensive applications (a performance hit of 10 percent will occur)
  - Applications for which cost is a primary issue

## RAIDset Planning Considerations

RAIDsets (RAID 3/5) are enhanced stripesets—they use striping to increase I/O performance and distributed-parity data to ensure data availability. [Figure 28](#) shows an example of a RAIDset that uses five members.



**Figure 28: Five-member RAIDset using parity**

RAIDsets are similar to stripesets in that the I/O requests are broken into smaller “chunks” and striped across the disk drives. RAIDsets also create chunks of parity data and stripe them across all the members of the RAIDset. Parity data is derived mathematically from the I/O data and enables the controller to reconstruct the I/O data if a single disk drive fails. Thus, it becomes possible to lose a disk drive without losing access to the data it contained. Data could be lost if a second disk drive fails before the controller replaces the first failed disk drive and reconstructs the data.

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rates. You can set the chunk size or use the default setting. See “[Chunk Size](#)” on page 89, for information about setting the chunk size.

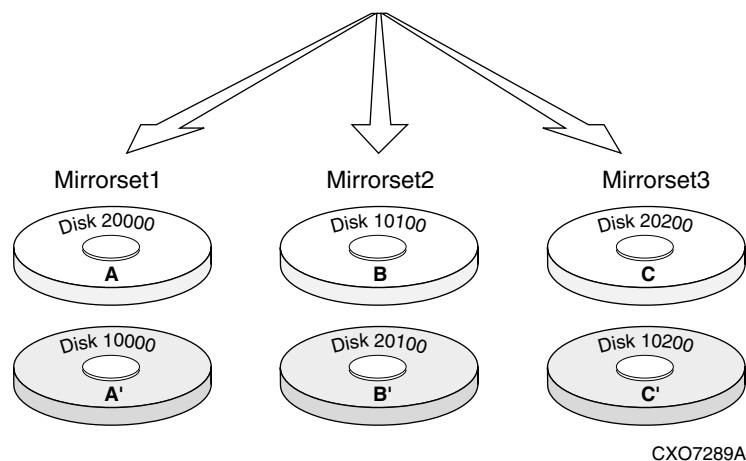
### Keep these points in mind when planning RAIDsets

- Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets.
- Both cache modules must be the same size.

- A RAIDset must include at least 3 disk drives, but no more than 14.
- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you waste 5 GB of capacity on each 9 GB member.
- RAIDsets are particularly well-suited for the following:
  - Small to medium I/O requests
  - Applications requiring high availability
  - High read request rates
  - Inquiry-type transaction processing
- RAIDsets are not particularly well-suited for the following:
  - Write-intensive applications
  - Database applications in which fields are continually updated
  - Transaction processing

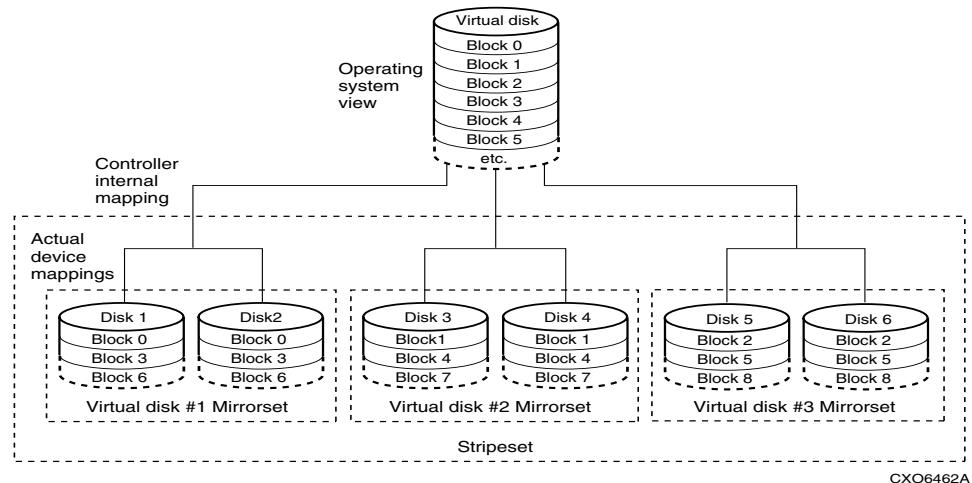
## Striped Mirrorset Planning Considerations

Striped mirrorsets (RAID 0+1) are a configuration of stripesets whose members are also mirrorsets (Figure 29). Consequently, this kind of storageset combines the performance of striping with the reliability of mirroring. The result is a storageset with very high I/O performance and high data availability. Figure 30 shows a second example of a striped mirrorset using six members.



**Figure 29: Striped mirrorset (example 1)**

The failure of a single disk drive has no effect on the ability of the storage set to deliver data to the host. Under normal circumstances, a single disk drive failure has very little effect on performance. Because striped mirrorsets do not require any more disk drives than mirrorsets, this storage set is an excellent choice for data that warrants mirroring.



**Figure 30: Striped mirrorset (example 2)**

Plan the mirrorset members, and plan the stripeset that will contain them. Review the recommendations in “[Planning Considerations for StorageSet](#)” on page 76, and “[Mirrorset Planning Considerations](#)” on page 80.

## StorageSet Expansion Considerations

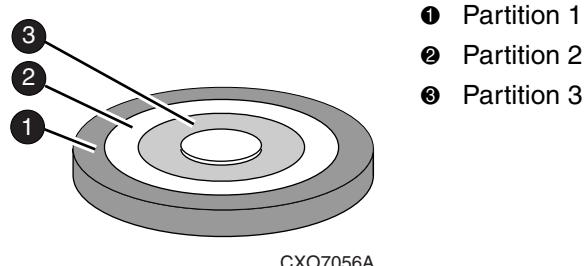
StorageSet Expansion allows for the joining of two of the same kind of storage containers by concatenating RAIDsets, stripesets, or individual disks, thereby forming a larger virtual disk, which is presented as a single unit. The *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* describes the CLI command, ADD CONCATSETS, which is used to perform concatenation.



**Caution:** Use the ADD CONCATSETS command only with host operating systems that support dynamic volume expansion. Use of this command could result in inaccessible data, if the operating system cannot handle one of its disks increasing in size.

## Partition Planning Considerations

Use partitions to divide a container (storageSet or individual disk drive) into smaller pieces, each of which can be presented to the host as its own storage unit. [Figure 31](#) shows the conceptual effects of partitioning a single-disk container.



**Figure 31: One example of a partitioned single-disk unit**

You can create up to eight partitions per storageSet (disk drive, RAIDset, mirrorset, stripeset, or striped mirrorset). Each partition has its own unit number so that the host can send I/O requests to the partition just as it would to any

unpartitioned storageset or device. Partitions are separately addressable storage units; therefore, you can partition a single storageset to service more than one user group or application.

## Defining a Partition

Partitions are expressed as a percentage of the storageset or single disk unit that contains them:

- Mirrorsets and single disk units—the controller allocates the largest whole number of blocks that are equal to or less than the percentage you specify.
- RAIDsets and stripesets—the controller allocates the largest whole number of stripes that are less than or equal to the percentage you specify.
  - Stripesets—the stripe size = chunk size  $\times$  number of members.
  - RAIDsets—the stripe size = chunk size  $\times$  (number of members minus 1)

An unpartitioned storage unit has more capacity than a partition that uses the whole unit because each partition requires a small amount of disk space for metadata.

## Guidelines for Partitioning Storagesets and Disk Drives

Keep these points in mind when planning partitions for storagesets and disks:

- Each storageset or disk drive may have up to eight partitions.
- In transparent failover mode, all partitions of a particular container must be on the same host port. Partitions cannot be split across host ports.
- In multiple-bus failover mode, all the partitions of a particular container must be on the same controller. Partitions cannot be split across controllers.
- Partitions cannot be combined into storagesets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.
- Just as with storagesets, you do not have to assign unit numbers to partitions until you are ready to use them.
- The CLONE utility cannot be used with partitioned mirrorsets or partitioned stripesets. (See “[Creating Clones for Backup](#)” on page 189 for details about cloning.)

# Changing Characteristics Through Switches

CLI command switches allow the user another level of command options. There are three types of switches that modify the storageset and unit characteristics:

- Storageset switches
- Initialization switches
- Unit switches

The following sections describe how to enable/modify switches. They also contain a description of the major CLI command switches.

## Enabling Switches

If you use SWCC to configure the device or storageset, you can set switches from SWCC during the configuration process, and SWCC automatically applies them to the storageset or device. See the SWCC online help for information about using SWCC.

If you use CLI commands to configure the storageset or device manually, the configuration procedure found in Chapter 5 of this guide will show you when and how to enable each switch. The *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* contains the details of the CLI commands and their switches.

## Changing Switches

You can change the RAIDset, mirrorset, device, and unit switches at any time. You cannot change the initialize switches without destroying data on the storageset or device. These switches are integral to the formatting and can only be changed by re-initializing the storageset.



**Caution:** Initializing a storageset is similar to formatting a disk drive; all data is destroyed during this procedure.

## Specifying StorageSet and Partition Switches

The characteristics of a particular storageSet can be set by specifying switches when the storageSet is added to the controllers' configuration. Once a storageSet has been added, the switches can be changed by using a SET command. Switches can be set for partitions and the following types of storageSets:

- RAIDset
- Mirrorset

Stripesets have no specific switches associated with their ADD and SET commands.

### RAIDset Switches

Use the following types of switches to control how a RAIDset ensures data availability:

- Replacement policy
- Reconstruction policy
- Remove/replace policy

For details on the use of these switches, refer to SET RAIDSET and SET *RAIDset-name* commands in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

### Mirrorset Switches

Use the following switches to control how a mirrorset behaves to ensure data availability:

- Replacement policy
- Copy speed
- Read source
- Membership

For details on the use of these switches, refer to ADD MIRRORSET and SET *mirrorset-name* commands in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

### Partition Switches

The following switches are available when creating a partition:

- Size
- Geometry

For details on the use of these switches, refer to **CREATE\_PARTITION** command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

## Specifying Initialization Switches

Initialization switches set characteristics for established storagesets before they are made into units. The following kinds of switches effect the format of a disk drive or storageset:

- Chunk Size (for stripesets and RAIDsets only)
- Save Configuration
- Destroy/Nodestroy
- Geometry

Each of these switches is described in the following sections.

---

**Note:** After initializing the storageset or disk drive, you cannot change these switches without reinitializing the storageset or disk drive.

---

### Chunk Size

With ACS software, a parameter for chunk size (chunksize=default or n) on some storagesets can be set. However, unit performance may be negatively impacted if a non-default value is selected as the chunksize.

If a non-default chunk size has been calculated, verify that the chunk size value is divisible by 8, with no remainder. If the value is not aligned with this rule, adjust the chunk size value upward until it divisible by 8, with no remainder.

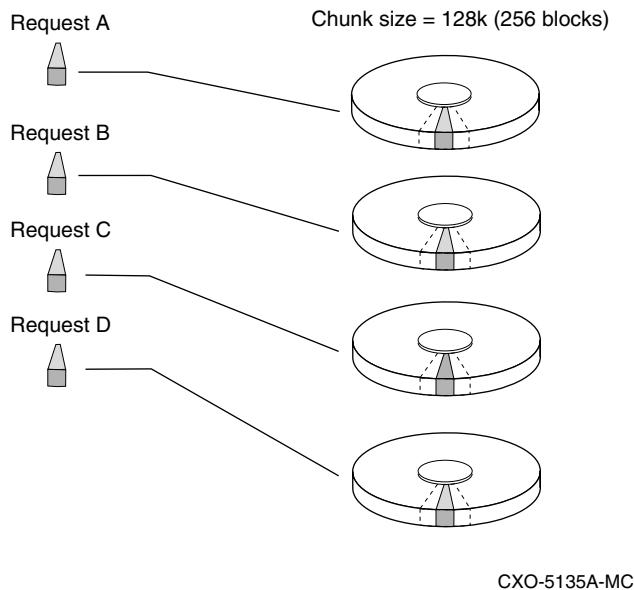
Specify the chunk size of the data to be stored to control the stripesize used in RAIDsets and stripesets:

- **CHUNKSIZE=DEFAULT** lets the controller set the chunk size based on the number of disk drives (d) in a stripeset or RAIDset. If number of drives is less or equal to 9, then chunk size = 256. If the number of drives is greater than 9, then chunk size = 128.

- `CHUNKSIZE=n` lets you specify a chunk size in blocks. The relationship between chunk size and request size determines whether striping increases the request rate or the data-transfer rate.

## Increasing the Request Rate

A large chunk size (relative to the average request size) increases the request rate by enabling multiple disk drives to respond to multiple requests. If one disk drive contains all of the data for one request, then the other disk drives in the storage set are available to handle other requests. Thus, separate I/O requests can be handled in parallel, which increases the request rate. This concept is shown in [Figure 32](#).



**Figure 32: Large chunk size increases request rate**

Large chunk sizes also tend to increase the performance of random reads and writes. HP recommends that you use a chunk size of 10 to 20 times the average request size, rounded to the closest prime number.

To calculate the chunk size that should be used for your subsystem, you must first analyze the types of requests that are being made to the subsystem:

- Many parallel I/Os that use a small area of disk should use a chunk size of 10 times the average transfer request rate.

- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate. If you do not know, then you should use a chunk size of 15 times the average transfer request rate.
- If you have mostly sequential reads or writes (like those needed to work with large graphic files), make the chunk size for RAID 0 and RAID 0+1 a small number (for example: 67 sectors). For RAID 5, make the chunk size a relatively large number (for example: 253 sectors).

Table 10 shows a few examples of chunk size selection.

**Table 10: Example Chunk Sizes**

Transfer Size (KB)	Small Area of I/O Transfers	Unknown	Random Areas of I/O Transfers
2	41	59	79
4	79	113	163
8	157	239	317

## Increasing Sequential Data Transfer Performance

RAID 0 and RAID 0+1 sets intended for high data transfer rates should use a relatively low chunk size (for example: 67 sectors). RAID 5 sets intended for high data rate performance should use a relatively large number (for example: 253 sectors).

## Save Configuration

The SAVE\_CONFIGURATION switch is for a single-controller configuration only. This switch reserves an area on each of the disks for the container being initialized. The controller can write subsystem configuration data on this area. If the controller is replaced, the new controller can read the subsystem configuration from the reserved areas of disks.

If you specify SAVE\_CONFIGURATION for a multi-device storage set, such as a stripeset, the complete subsystem configuration is periodically written on each disk in the storage set.



**Caution:** If user data already exists on a storageset, do NOT reinitialize it with the *save configuration* option, as this will change the site and position of the user data on the storageset. Always back up user data prior to reinitializing any storageset.

---

The SHOW DEVICES FULL command shows which disks are used to backup configuration information.

---

**Note:** DO NOT use SAVE\_CONFIGURATION in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

---

**Note:** HP recommends that you do not use SAVE\_CONFIGURATION on every unit and device on the controller.

---

## Destroy/Nodestroy

Specify whether to destroy or retain the user data and metadata when a disk is initialized after it has been used in a mirrorset or as a single-disk unit.

---

**Note:** The DESTROY and NODESTROY switches are only valid for mirrorsets and striped mirrorsets.

---

- DESTROY (default) overwrites the user data and forced-error metadata when a disk drive is initialized.
- NODESTROY preserves the user data and forced-error metadata when a disk drive is initialized. Use NODESTROY to create a single-disk unit from any disk drive that has been used as a member of a mirrorset. See the REDUCED command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for information on removing disk drives from a mirrorset.

NODESTROY is ignored for members of a RAIDset.

## Geometry

The geometry parameters of a storageset can be specified. The geometry switches are:

- CAPACITY—the number of logical blocks. The range is from 1 to the maximum container size.
- CYLINDERS—the number of cylinders used. The range is from 1 to 16777215.
- HEADS—the number of disk heads used. The range is from 1 to 255.
- SECTORS\_PER\_TRACK—the number of sectors per track used. The range is from 1 to 255.

## Specifying Unit Switches

Several switches control the characteristics of units. The unit switches are described under the *SET unit-number* command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

One unit switch, *ENABLE/DISABLE\_ACCESS\_PATH*, determines which host connections can access the unit, and it is described in the larger topic of matching units to specific hosts. This topic is described in Chapter 1, “[Determining Connections](#)” on page 39.

## Creating Storage Maps

Configuring a subsystem will be easier if you know how the storagesets, partitions, and JBODs correspond to the disk drives in your subsystem. You can more easily see this relationship by creating a hardcopy representation, also known as a storage map.

To make a storage map, fill out the templates provided in Appendix A as you add storagesets, partitions, and JBOD disks to the configuration and assign them unit numbers. Label each disk drive in the map with the higher levels it is associated with, up to the unit level.

## Using LOCATE Command to Find Devices

If you want to complete a storage map at a later time but do not remember where the disk drives and partitions are located, use the CLI command LOCATE. The LOCATE command flashes the (fault) LED on the drives associated with the specific storageset or unit. To turn off the flashing LEDs, enter the CLI command LOCATE CANCEL.

The following procedure is an example of the commands to locate all the disk drives that make up unit D104:

1. Enter the following command:

```
LOCATE D104
```

The LEDs on the disk drives that make up unit D104 will flash.

2. Note the position of all the drives contained within D104.
3. Enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```

The following procedure is an example command to locate all the drives that make up RAIDset R1:

1. Enter the following command:

```
LOCATE R1
```

2. Note the position of all the drives contained within R1.
3. Enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```

## Example Storage Map—Model 4310R Disk Enclosure

[Table 11](#) shows an example of four Model 4310R disk enclosures (single-bus I/O).

**Table 11: Model 4310 disk enclosure, example of storage map**

Model 4310R Disk Enclosure Shelf 4 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1	D101 S1 M2	D103 M4	D105	D107	D108 S3	D1 S4 M6	D2 R3	D3 S5	spare
Model 4310R Disk Enclosure Shelf 1 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1	D101 S1 M1	D102 M3	D104 S2	D106 R2	D108 S3	D1 S4 M5	D2 R3	D3 S5	D4 M7

**Table 11: Model 4310 disk enclosure, example of storage map (Continued)**

Model 4310R Disk Enclosure Shelf 2 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1	D101 S1 M1	D102 M3	D104 S2	D106 R2	D108 S3	D1 S4 M5	D2 R3	D3 S5	D4 M7
Model 4310R Disk Enclosure Shelf 3 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1	D101 S1 M2	D103 M4	D104 S2	D106 R2	D108 S3	D1 S4 M6	D2 R3	D3 S5	spare

The following notes describe the information contained in [Table 11](#):

- Unit D100 is a 4-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, and Disk40000.
- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
  - M1 is a 2-member mirrorset consisting of Disk10100 and Disk20100.
  - M2 is a 2-member mirrorset consisting of Disk30100 and Disk40100.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10200 and Disk20200.

- Unit D103 is a 2-member mirrorset named M4. M4 consists of Disk30200 and Disk40200.
- Unit D104 is 3-member stripeset named S2. S2 consists of Disk10300, Disk20300, and Disk30300.
- Unit D105 is a single (JBOD) disk named Disk40300.
- Unit D106 is a 3-member RAID 3/5 storageset named R2. R2 consists of Disk10400, Disk20400, and Disk30400.
- Unit D107 is a single (JBOD) disk named Disk40400.
- Unit D108 is a 4-member stripeset named S3. S3 consists of Disk10500, Disk20500, Disk30500, and Disk40500.
- Unit D1 is a 2-member striped mirrorset named S4. S4 consists of M4 and M5:
  - M5 is a 2-member mirrorset consisting of Disk10800 and Disk20800.
  - M6 is a 2-member mirrorset consisting of Disk30800 and Disk40800.
- Unit D2 is a 4-member RAID 3/5 storageset named R3. R3 consists of Disk11000, Disk21000, Disk31000, and Disk41000.
- Unit D3 is a 4-member stripeset named S5. S5 consists of Disk11100, Disk21100, Disk31100, and Disk41100.
- Unit D4 is a 2-member mirrorset named M7. M7 consists of Disk11200 and Disk21200.
- Disk31200 and Disk41200 are spareset members.



# Preparing the Host System

3

This chapter describes how to prepare your Tru64 UNIX host computer to accommodate the HSG80 controller storage subsystem.

The following information is included in this chapter:

- [Installing RAID Array Storage System](#), page 100
- [Making a Physical Connection](#), page 104
- [Verifying Disk Array Device Recognition](#), page 105
- [DECsafe Available Server Environment \(ASE\)](#), page 112
- [HSG80 Units and Tru64 UNIX Utilities](#), page 112
- [Solution Software Upgrade Procedures](#), page 113
- [New Features, ACS 8.8-1 for Tru64](#), page 115

Refer to Chapter 4 for instructions on how to install and configure the HSG Agent. The Agent for HSG is operating system-specific and polls the storage.

## Installing RAID Array Storage System

---



**WARNING:** A shock hazard exists at the backplane when the controller enclosure bays or cache module bays are empty.

Be sure the enclosures are empty, then mount the enclosures into the rack. DO NOT use the disk enclosure handles to lift the enclosure. The handles cannot support the weight of the enclosure. Only use these handles to position the enclosure in the mounting brackets.

Use two people to lift, align, and install any enclosure into a rack. Failure to use two people might cause personal injury and/or equipment damage.

---



**Caution:** Controller and disk enclosures have no power switches. Make sure the controller enclosures and disk enclosures are physically configured before turning the PDU on and connecting the power cords. Failure to do so can cause equipment damage.

---

1. Be sure the enclosures are empty before mounting them into the rack. If necessary, remove the following elements from the controller enclosure:
  - Environmental Monitoring Unit (EMU)
  - Power Supplies
  - External Cache Batteries (ECBs)
  - Fans

If necessary, remove the following elements from the disk enclosure:

- Power Supply/Blower Assemblies
- Disk Drives
- Environmental Monitoring Unit (EMU)
- I/O Modules

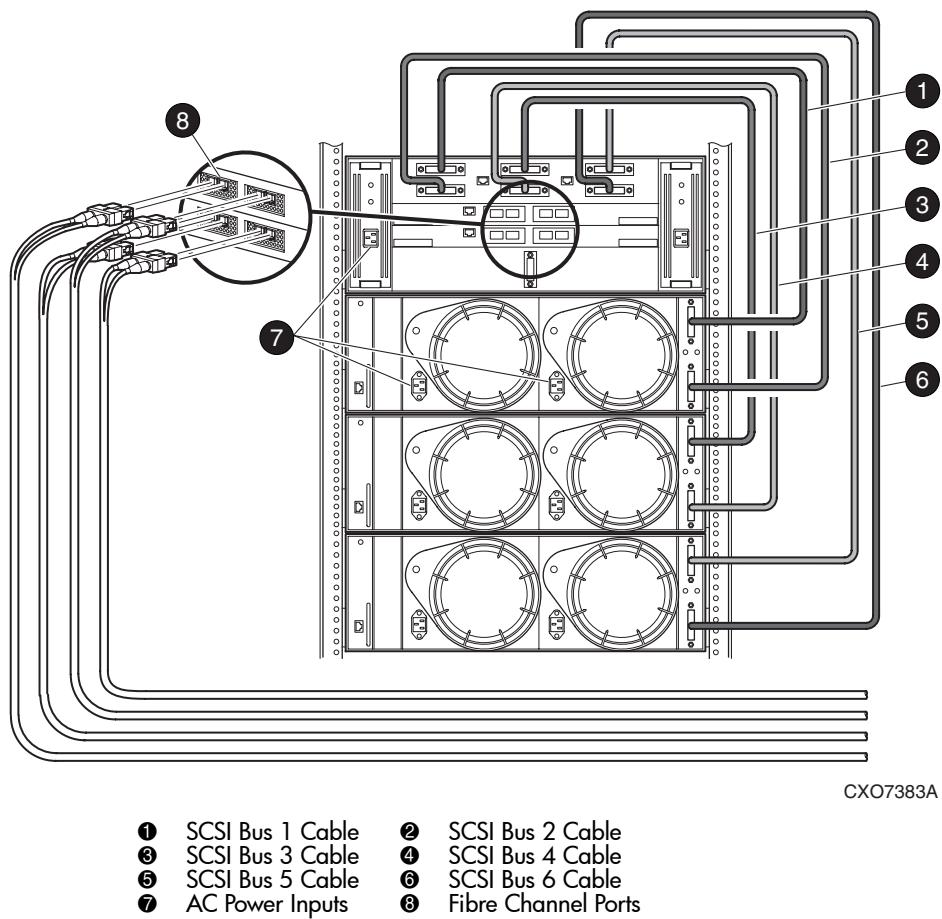
Refer to the *StorageWorks Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide* and the *StorageWorks Model 4300 Family Ultra3 LVD Disk Enclosures User Guide* for further information.

2. Install brackets onto the controller enclosure and disk enclosures. Using two people, mount the enclosures into the rack. Refer to the mounting kit documentation for further information.

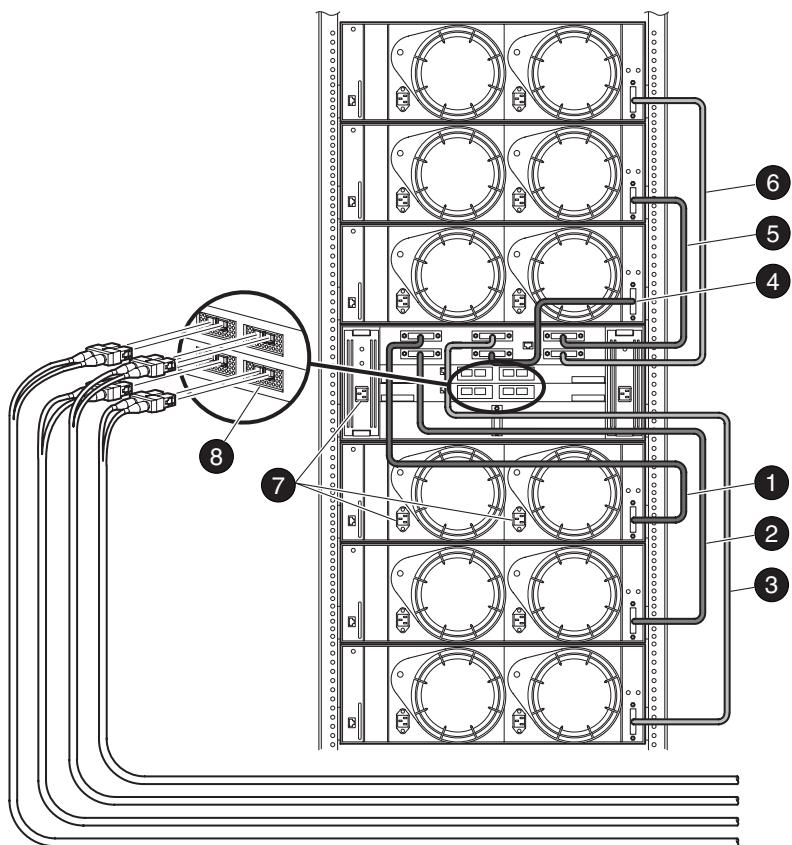
3. Install the elements. Install the disk drives. Make sure you install blank panels in any unused bays.

Fibre Channel cabling information is shown to illustrate supported configurations. In a dual-bus disk enclosure configuration, disk enclosures 1, 2, and 3 are stacked below the controller enclosure—two SCSI Buses per enclosure (see [Figure 33](#)). In a single-bus disk enclosure configuration, disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure—one SCSI Bus per enclosure (see [Figure 34](#)).

4. Connect the six VHDCI UltraSCSI bus cables between the controller and disk enclosures, as shown in [Figure 33](#) for a dual-bus system and [Figure 34](#) for a single-bus system. Note that the supported cable lengths are 1, 2, 3, 5, and 10 meters.
5. Connect the AC power cords from the appropriate rack AC outlets to the controller and disk enclosures.



**Figure 33: Dual-Bus Enterprise Storage RAID Array Storage System**



CXO7382A

①	SCSI Bus 1 Cable	②	SCSI Bus 2 Cable
③	SCSI Bus 3 Cable	④	SCSI Bus 4 Cable
⑤	SCSI Bus 5 Cable	⑥	SCSI Bus 6 Cable
⑦	AC Power Inputs	⑧	Fibre Channel Ports

**Figure 34: Single-Bus Enterprise Storage RAID Array Storage System**

## Making a Physical Connection

To attach a host computer to the storage subsystem, install one or more host bus adapters into the computer. A Fibre Channel (FC) cable goes from the host bus adapter to an FC switch.

### Preparing to Install Host Bus Adapter

Before installing the host bus adapter:

1. Perform a complete backup of the entire system.
2. Shut down the computer system or perform a hot addition of the adapter based upon directions for that server.

### Installing Host Bus Adapter

To make a physical connection, first install a host bus adapter.



**Caution:** Protect the host bus adapter board from electrostatic discharge by wearing an ESD wrist strap. DO NOT remove the board from the antistatic cover until you are ready to install it.

---

You need the following items to begin:

- Host bus adapter board
- The computer hardware manual
- Appropriate tools to service your computer

The host bus adapter board plugs into a standard PCI slot in the host computer. Refer to the system manual for instructions on installing PCI devices.

---

**Note:** Take note of the worldwide name (WWN) of each adapter. Do not power on anything yet. For the FC switches to autoconfigure, power on equipment in a certain sequence. Also, the controllers in the subsystem are not yet configured for compatibility with Tru64 UNIX.

---

## Verifying Disk Array Device Recognition

Use the `scu show edt` command at the UNIX prompt to see the list of new disk array devices. The device files are created for each logical unit automatically in Tru64 UNIX during system startup.

1. Verify the character-type device files have been automatically created:

Tru64 v4 example:

```
# file /dev/rrz*
```

Tru64 v5 example:

```
# file /dev/rdisk/dsk*
```

2. Verify the block-type device files have been automatically created:

Tru64 v4 example:

```
# file /dev/rz*
```

Tru64 v5 example:

```
# file /dev/disk/dsk*
```

You can verify the allocation for the device files by calculating as follows:

Block type: `rzXYZ`

Character type: `rrzXYZ`

X = b through h = LUN1 through LUN7 (no letter is used for LUN0)

Y = SCSI bus number  $\times$  8 + SCSI ID

Z = partition (a through h)

---

**Note:** Note: Target ID > 8 is not allowed for hosts.

---

### Example of calculation

In this example, OPEN-3 of LUNs 0 through 7 with SCSI ID = 0 is connected to the card containing SCSI bus number = 1. With this calculation,  $Y=1\times8+0=8$ , the following device files are allocated to each partition of LUN1 (X = b):

```
rzb8a,rrzb8a,rzb8b,rrzb8b,rzb8c,rrzb8c,  
rzb8d,rrzb8d,rzb8e,rrzb8e,rzb8f,rrzb8f,  
rzb8g,rrzb8g,rzb8h,rrzb8h
```

## Configure Disk Devices

Configure the disk array devices in the same way you would configure any new disk on the host. Creating scripts to configure all devices at once may save you considerable time.

Configuring devices typically requires these steps:

1. Partition and label each device using the `disklabel` command.
2. Create a file system for each device using the `newfs` command.
3. Create a mount directory for each device using the `mkdir` command.
4. Enter each device into the mount table by editing `/etc/fstab`.
5. Reboot the host. After bootup use the `df` command to verify the devices auto-mounted.
6. Verify file system operation by copying a file to each device and then deleting it.

## Writing the Partition Labels

Use the `disklabel` command to label the partition for each logical unit. Partition c specifies the entire area in the logical unit. Check that no errors are found in the partition settings after the labeling.

You can edit the disk partition size using the `disklabel` command with option `-e`. When the `disklabel -e` command is executed, the `vi` editor for the environment in which you are working starts up. After completing the editing, save the file and execute the `disklabel` command again; the partition setting is renewed.

To write the partition label:

1. Enter `disklabel -rw`.

Tru64 v4

```
# disklabel -rw rz8 HSG80
```

2. Specify the device file in the `rzXY` format (X = LUN, Y = SCSI bus number  $\times$  8 + SCSI ID).
3. Specify the disk name that is registered for that device file.

Tru64 v5

```
# disklabel -rw dsk10 HSG80
```

4. Enter `disklabel -r` to verify labeling and partition settings.

Tru64 v4

```
# disklabel -r rz8
```

5. Specify the device file in the `rzXY` format (X = LUN, Y = SCSI bus number  $\times 8 +$  SCSI ID).

Tru64 v5

```
# disklabel -r dsk10
```

## Creating the File Systems

Create a file system for each new OPEN-x device. Optionally, you can create and use an advanced file system (AdvFS) to overcome the size and speed limitations of the file system. If you are not sure which file system is right for your setup, contact HP customer support.

### Creating a File System For Each Device

To create a file system for each device, simply type `# newfs device_file_name`

### Creating an Advanced File System

Creating an advanced file system can mean creating a new advanced file system domain, or you adding a new fileset to an existing domain.

If you allocate multiple disk partitions to a domain, the advanced file system utilities must be installed. Installing the utilities will enable you to allocate multiple disk partitions to a domain using the `addvol` command, for example:

```
# addvol /dev/rzb8c domain1
```

#### To create a new domain:

```
# mkgfdmn device_file_name domain_name
```

Tru64 v4 Example

```
# mkgfdmn /dev/rz12c domain1.
```

The example uses the block-type device file name (rzXYZ) to create an advanced file system domain for device rz12c.

Tru64 v5 Example

```
# newfs /dev/rdisk/dsk10c HSG80
```

To create a new `fileset` in the new or existing advanced file system domain:

```
# mkfset domain_name fileset_name
```

For example, to create `fileset1` in `domain1`:

```
# mkfset domain1 fileset1
```

## Creating Mount Directories

Create a mount directory for each device and assign each mount directory a unique name that identifies the device being mounted.

1. Create a mount directory:

```
# mkdir /mount_directory_name
```

For example, to create a mount directory for LUN 2 (partition `c`) on the disk array, enter `# mkdir /HSG80`.

2. Verify the new mount directory:

```
# ls /
```

## Mounting the File Systems

After the file systems and mount directories have been created, you can mount the file system for each new device.

### To mount a Tru64 file system

1. Mount device:

```
# mount device_file_name mount_directory
```

For example, to mount device `rz12c` with mount directory name `HSG80`, enter:

```
# mount /dev/rz12c /HSG80.
```

For Tru64 v5:

```
# mount /dev/disk/dsk10c /mnt
```

2. Assign the appropriate ownership and permissions:

```
# chown owner:group *device_file_name*
```

For example, to assign ownership to `rz12c` with owner Oracle, group dba, enter:

```
# chown oracle:dba *rz12c*.
```

### To mount an advanced file system

1. Mount the file system:

```
# mount -t advfs domain_name#fileset_name mount_directory
```

For example, to mount the file system with mount directory HSG80, enter:

```
mount -t advfs domain1#fileset1/HSG80.
```

2. Assign the appropriate ownership and permissions:

```
# chown owner:group *device_file_name*
```

For example, to assign ownership to rz12c with owner Oracle, group dba, enter:

```
# chown oracle:dba *rz12c*.
```

## Verifying the File Systems

Verify that the new file systems were created correctly and are functioning properly.

1. Display all mounted file systems by entering:

```
# df
```

The default display for drive capacity is 512-byte blocks. To view the capacity in KB rather than in 512-byte blocks, enter:

```
# df -k
```

2. Go to a new device directory:

```
# cd /mount_directory
```

For example, enter:

```
cd /HSG80.
```

3. Copy a file from the root directory to the new device:

```
# cp /filename file_name.back1
```

For example, to copy file vmunix from the root directory to the HSG80 device, enter:

```
cp /vmunix vmunix.back1.
```

4. Copy a file to the new device again:

```
# cp /filename file_name.back2
```

For example, to copy the same file again, enter:

```
cp /vmunix vmunix.back2.
```

5. List the files in the current directory:

```
# ls -l
```

The vmunix.back1 and vmunix.back2 files should be shown.

6. Delete the files you copied:

```
# rm file_name
```

For example, enter `rm vmunix.back1` to remove the file `vmunix.back1`.

## Setting and Verifying Automatic Mounting at Bootup

The `/etc/fstab` file contains boot time mounting parameters for disk devices. To set and verify automatic mounting at bootup:

1. Edit the `/etc/fstab` file after first making a backup copy of the file. Add a line for each new device to be mounted.

*Example*

```
#vi /etc/fstab
```

**Table 12: Boot time mounting parameters**

Parameter 1	Parameter 2	Parameter 3	Parameter 4	Parameter 5	Parameter 6
<code>/dev/rz4a</code>	/	ufs	rw	1	1
<code>/proc</code>	<code>/proc</code>	procfs	rw	0	0
<code>/dev/rz12a</code>	<code>/usr</code>	ufs	rw	1	2
<code>/dev/rz12b</code>	<code>swap1</code>	ufs	rw	0	2
<code>/dev/rz12</code>	<code>/HSG80</code>	ufs	rw	1	3

**Table 13: Mounting parameter descriptions**

Parameter	Description	Enter
1	Device to mount	Block-type device file name
2	Mount point	Mount directory name
3	File system	Type of file system (for example, UFS, AdvFS)

**Table 13: Mounting parameter descriptions (Continued)**

Parameter	Description	Enter
<b>4</b>	Mount options	Options (that is, <code>rw</code> for "read/write")
<b>5</b>	Frequency	The <code>dump</code> command controls file system backup. Typing "1" backs up the file system. Typing "0" prevents backup.
<b>6</b>	<code>fsck</code> order (File System Check—For UFS systems only, you must <code>fsck</code> disks that contain mount points before mounting other disks on those mount points.)	Block-type device file name.

2. Shut down and reboot the system.
3. Use the `df` or `df -k` command to verify file system auto mounting.

## Checking path failover

The disk array supports Tru64 path failover (Tru64 v5 only). You can connect multiple HBAs to the disk array with shared LUNs. Confirm the existence of multiple paths for devices as follows:

1. Type `hwmgr -view device` to obtain the HWID for the device.
2. Type `hwmgr -show scsi -full -id HWID` to confirm the status of paths to the device.

If more than one path is currently connected, the status of each path shows as `Valid`. If you change the cabling configuration, the old paths will show as `Stale`. Use the `hwmgr -refresh` command to remove the stale paths.

---

**Caution:** Tru64 version 5 fully supports dynamic load balancing and failover when multiple HBAs are connected to the same LUN. This feature is **ONLY** available with version 5.

---

On Tru64 version 4, DO NOT connect multiple HBAs to the same LUN and DO NOT allow the same LUN to be accessible from multiple HBAs on Tru64 version 4 hosts. ON the version 4 system, this will allow multiple device names to access the same LUN. If both device names are used, data corruption will result.

---

## DECsafe Available Server Environment (ASE)

HSG80 disk devices can be used with the DECsafe Available Server Environment (ASE) for Tru64 UNIX provided a valid Host configuration (including Host adapters) is used to support them. Refer to the *Tru64 UNIX ASE Installation and User's Guide, Software Product Description, SPD: 44.17.xx*, for further information. Refer to the *Release Notes* for supported host adapters and Tru64 UNIX version levels for ASE.

## Using genvmunix

---

**Note:** This section on *genvmunix* does not apply to Tru64 UNIX V5.1x, since the configuration file has a different format for devices.

---

If you use *genvmunix* to initialize the system and *doconfig* to build a new configuration file, the new configuration file will only list the HSG80 LUN 0 units; nonzero LUNs will not appear in the new configuration file.

Before rebuilding a configuration file using *genvmunix*, save any existing customized configuration file that has entries for HSG80 units. After rebuilding the configuration file, add the entries from the saved configuration file to the new configuration file.

## HSG80 Units and Tru64 UNIX Utilities

This section contains notes on the interactions of some Tru64 UNIX utilities with storage units in your Enterprise Storage RAID Array storage system.

### Add recognition for Fibre Channel adapters

Make sure you have the latest patch kit installed for Tru64. Patch kits are available at the following web site:

<http://tru64unix.compaq.com/support.html>

## Solution Software Upgrade Procedures

Use the following procedures for upgrades to your Solution Software. It is considered best practice to follow this order of procedures:

1. Perform backups of data prior to upgrade.
2. Verify operating system versions, upgrade operating systems to supported versions and patch levels.
3. Quiesce all I/O and unmount all file systems before proceeding.
4. Upgrade solution software.
5. Upgrade ACS software.

---

**Note:** Solely for the purpose of performing upgrades to the ACS firmware, this Solution Software Kit supports previous ACS V8.6 and V8.7. Do not mix ACS versions in the same SAN.

---

Refer to the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide* and the *Solution Software Release Notes* for the latest information on upgrades.

Use the following installation and upgrade procedures for Tru64 UNIX systems:

---

**Note:** You must have the SNMP service installed on your client computer before the installation.

---

1. Insert the Solution Software CD-ROM into a computer running Windows 2000 or Windows NT 4.0 with Service Pack 4 or later.
2. Using Microsoft Windows Explorer, go to the SWCC directory on the CD-ROM and double-click `setup.exe`. The SWCC Setup window will appear.
3. Select **HSG80 Controller for ACS85 new** and click **Next**. Follow the instructions on the screen.

### Installing Agent

1. Insert the Solution Software CD-ROM into the host computer.

2. Type one of the following commands at the command prompt, depending on your operating system version.

For Tru64 UNIX Version 4.0x, type:

```
# mount -r -t cdfs -o rrip /dev/rz6c /mnt
```

Substitute for `rz6c`, if necessary, for your CD-ROM.

For Tru64 UNIX Version 5.1x, type:

```
# mount -r -t cdfs -o rrip /dev/disk/cdrom0c /mnt
```

Substitute for `cdrom0c`, if necessary, for your CD-ROM.

3. Change directories on the CD-ROM by typing:

```
# cd /mnt/agent
```

4. To run the installation program, type the following at the command prompt:

```
# setld -l . ( -l is a lowercase L)
```

You are asked if you want to install the listed subsets.

5. Enter the corresponding number of the software and then press **RETURN**.

6. Select option **y**, and then press **RETURN**.

7. Follow the instructions on the screen.

## Upgrading the Agent and ACS on Standalone Servers

1. For any Tru64 UNIX version supported, stop the Agent by using the following steps:

- a. Stop I/O to your FC drives.

- b. Unmount your FC drives using the `umount` command.

- c. Stop the Agent by typing the following commands:

```
# cd /usr/opt/SWCCXXX/scripts
```

```
# execute swcc_config
```

- d. From the SWCC Agent Configuration Utility choose **7**) to disable the Agent (`steamd`)

2. To delete the Agent configuration files, type the following:

```
# setld -d SWCCXXX
```

At the prompt:

Answer **Y** - Your configuration files will be deleted.

Answer **N** - The old configuration files will be kept.

3. To install the new Agent, follow the steps listed in "Installing Agent" above.

---

**Note:** If in a Cluster, for Version 8.8-1, the Agent may be controlled by caa daemon.

---

4. Stop the SWCC Client from running on your Windows Client machine.
  - a. Remove the SWCC Client from the Windows Client machine using:  
Control Panel > Add/Remove Programs
  - b. Remove all entries associated with SWCC.
5. To install the new Client, follow the steps listed in "Installing Agent" above.
6. Upgrade your ACS to Version 8.8-1.
7. Remount your FC drives or reconfigure your storage as required.

To rescan the bus on Tru64 Version 5.x systems, type the following:

```
# hwmgr -scan scsi  
# dsfmgr -K  
# hwmgr -v dev
```

To rescan the bus on Tru64 Version 4.x systems, type the following:

```
# scu scan edt  
# scu show edt
```

8. Restart I/O (applications).
9. Start Client.

## New Features, ACS 8.8-1 for Tru64

The following are new features implemented in ACS 8.8-1 for use by the Tru64 operating system:

- [Host Connection Table Management Improvements](#), page 116
- [Selective Management Presentation](#), page 120
- [Linking WWIDs for Snap and Clone Units](#), page 123
- [SMART Error Eject](#), page 125
- [Error Threshold for Drives](#), page 129

## Host Connection Table Management Improvements

Current implementation of host connectivity grants access to the first ninety-six (96) host connections that attempt to login to the controller. After that, all edits to the host connection table are manual and require extensive CLI commands to delete and/or replace connections.

### Host Connection Table Locking

Host table lock and unlock commands have been added to control the connection table in NVRAM. When the table is locked, the host login request (PLOGI) will be rejected (unless the connection is already in the table) and the request will be stored internally on a rejected hosts table. In the default mode, if a PLOGI is received when the connection table is unlocked, the connection will be granted if there is room in the connection table. The lock state can be changed with the following CLI commands:

```
CLI> SET <THIS | OTHER> CONNECTIONS_LOCKED  
CLI> SET <THIS | OTHER> CONNECTIONS_UNLOCKED
```

The CONNECTIONS\_LOCKED and CONNECTIONS\_UNLOCKED switches must be typed completely to stop someone from inadvertently changing the state of the lock.

The lock is maintained in the failover information (fi) section of each controller's NV. When the state of the lock is changed on one controller, the other controller is updated as well. The existing CLI command to ADD CONN is not affected by the state of the lock.

### Viewing Host Connection Table Lock State

The state of the lock can be displayed using:

```
CLI> SHOW <THIS | OTHER>
```

The following string is displayed just before the port topology information:

```
Host Connection Table is <LOCKED | NOT locked >
```

## Example of Host Connection Table Unlock: (new output shown in bold)

```
AP_Bot> show this
Controller:
  HSG80      (C) DEC CX00000001 Software V87
Hardware 0000
  NODE_ID      = 5000-1FE1-FF00-0090
  ALLOCATION_CLASS = 1
  SCSI_VERSION    = SCSI-3
  Configured for dual-redundancy with ZG02804912
  In dual-redundant configuration
  Device Port SCSI address 6
  Time: 10-SEP-2001 15:45:54
  Command Console LUN is lun 0 (IDENTIFIER = 99)
  Host Connection Table is NOT locked
Host PORT_1:
  Reported PORT_ID = 5000-1FE1-FF00-0091
  PORT_1_TOPOLOGY = FABRIC (standby)
Host PORT_2:
  Reported PORT_ID = 5000-1FE1-FF00-0092
  PORT_2_TOPOLOGY = FABRIC (fabric up)
  Address        = 011200
  NOREMOTE_COPY
Cache:
  512 megabyte write cache, version 0022
  Cache is GOOD
  No unflushed data in cache
  CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
Mirrored Cache:
  Not enabled
Battery:
  NOUPS
  FULLY CHARGED
  Expires:        07-AUG-2003
```

Example of Host Connection Table Locked: (new output shown in bold)

```
AP_Bot> show this
Controller:
  HSG80      (C) DEC CX00000001 Software XC21P-0,
  Hardware 0000
    NODE_ID          = 5000-1FE1-FF00-0090
    ALLOCATION_CLASS = 1
    SCSI_VERSION     = SCSI-3
    Configured for dual-redundancy with ZG02804912
      In dual-redundant configuration
    Device Port SCSI address 6
    Time: 10-SEP-2001 15:48:24
    Command Console LUN is lun 0 (IDENTIFIER = 99)
    Host Connection Table is LOCKED

  Host PORT_1:
    Reported PORT_ID = 5000-1FE1-FF00-0091
    PORT_1_TOPOLOGY = FABRIC (standby)

  Host PORT_2:
    Reported PORT_ID = 5000-1FE1-FF00-0092
    PORT_2_TOPOLOGY = FABRIC (fabric up)
    Address          = 011200
    NOREMOTE_COPY

  Cache:
    512 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
    CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

  Mirrored Cache:
    Not enabled

  Battery:
    NOUPS
    FULLY CHARGED
    Expires:          07-AUG-2003
```

The state of the connection can be displayed using:

```
CLI> SHOW CONN <FULL>
```

<<< LOCKED >>> appears in the title area when the connection table is locked. If unlocked, or not supported (HOST\_FC only), the title area looks the same as it did for ACS version 8.7. The full switch displays the rejected hosts, with an index.

## Adding Rejected Host Connections to Locked Host Connection Table

With ACS version 8.8-1, it is now possible to keep track of rejected hosts in a list, and to keep this list synchronized across controllers. An index is now added to the record to aid the user in manually adding rejected connections. The command to manually add rejection connections is:

```
ADD CONNECTION REJECTED_HOST <index>
```

This adds the connection to the connection table in an OFFLINE state. The host must issue a FC plogi to make the connection active. There are mechanisms to do this in UNIX, VMS, but not NT (except reboot).

One way to force the connection into an online state is to do the following:

```
CLI> SET <THIS | OTHER> PORT_<1|2>_TOPOLOGY = OFFLINE, followed  
by
```

```
CLI> SET <THIS | OTHER> PORT_<1|2>_TOPOLOGY = FABRIC
```

This will force all hosts connected to that controller/port to login again. At the same time, hosts connected to the controller pair through the same switch (regardless of controller/port) will need to login again

---

**Note:** This implementation is the safest, since pinging the fabric name server would result in all hosts logging in again (up to 96).

---

When the connection is added, it gets deleted from the reject list. Index numbers for remaining rejected hosts are re-ordered.

## Implementation Notes

- For an upgrade to version 8.8-1 - The connection table is unlocked by default.
- For Backward Compatibility - Lock the table, and everything work the same as version 8.7.

- To Create a new SAN - Basically the system administrator unlocks the connection table, connect the desired hosts, and then lock the connection table. As the hosts are connected they login to the controller pair. After the connection table is locked, the host logins are rejected until the system administrator manually adds the host to the connection table.
- To Add a new Host to a SAN - A new host is added to the fabric that needs connectivity to the HSG80. Attempts to login are rejected because the connection table is locked. The system administrator is called, and manually adds an entry for the new host by creating a new connection from the rejected host.
- To Delete a Host - While the connection table is locked, delete the connection for the selected host. When the host realizes that it cannot talk to the HSG80 anymore, it may try to login again, but those attempts will be rejected since the connection table is locked.
- Too Many Hosts? - If more than 96 connections are present, all host logins will be rejected, regardless of the state of the lock.

## Selective Management Presentation

Selective Management Presentation is a control method that extends the use of the Selective Storage Presentation concept currently available on logical units (LUNs). You can use this mechanism to send control commands to the HSG80 controller (see “[What is Selective Storage Presentation?](#)” on page 48).

The new access mechanism provides control over which, SAN Management Agent systems can perform management operations.

In order to define the set of Management Agent host systems that can access the HSG80 management functions, new CLI commands are defined. These commands allow the user to selectively enable or disable host access to the control mechanism. These new commands provide for the addition and removal of Management Agent Host systems, and the ability to display the currently enabled systems.

## Removing Management Agent Host systems

The following command disables access to the management functions. The user can specify all systems or a list of systems.

```
HSG80> SET DISABLE_MANAGERS=ALL  
HSG80> SET DISABLE_MANAGERS=(host list...)
```

## Adding Management Agent Host systems

The following command enables access to the management functions. The user can specify all systems, or a list of systems.

```
HSG80> SET ENABLE_MANAGERS=ALL
- or -
HS80> SET ENABLE_MANAGERS=(host list...)
```

## Display Enabled Management Agents

The following command displays a list of the systems currently enabled to perform management functions.

```
HSG80> SHOW MANAGERS
Connection
      Name      Operating System      Controller Port      Address
      Status
      NEWCON14      AIX            THIS            1      011000      OL
      this          0
      HOST_ID=2000-0000-C922-46E2
      ADAPTER_ID=1000-0000-C922-46E2
      NEWCON15      WINNT           THIS            1      011200      OL
      this          0
      HOST_ID=2000-0000-C927-6735
      ADAPTER_ID=1000-0000-C927-6735
      NEWCON16      AIX            OTHER           1      011100      OL
      other          0
      HOST_ID=2000-0000-C925-0096
      ADAPTER_ID=1000-0000-C925-0096
      NEWCON17      WINNT           OTHER           1      011300      OL
      other          0
      HOST_ID=2000-0000-C923-2CD2
      ADAPTER_ID=1000-0000-C923-2CD2
```

In the event that all connections are enabled the display appears as follows.

```
HSG80> SHOW MANAGERS
      Connection          <<<All Connections Enabled>>>
      Name      Operating System   Controller Port   Address
      Status
NEWCON14      AIX           THIS             1       011000   OL
this          0
      HOST_ID=2000-0000-C922-46E2
      ADAPTER_ID=1000-0000-C922-46E2
NEWCON15      WINNT          THIS             1       011200   OL
this          0
      HOST_ID=2000-0000-C927-6735
      ADAPTER_ID=1000-0000-C927-6735
NEWCON16      AIX           OTHER            1       011100   OL
other         0
      HOST_ID=2000-0000-C925-0096
      ADAPTER_ID=1000-0000-C925-0096
NEWCON17      WINNT          OTHER            1       011300   OL
other         0
      HOST_ID=2000-0000-C923-2CD2
      ADAPTER_ID=1000-0000-C923-2CD2
      <<<All Connections Enabled>>>
```

## Enabling SAN Security

To enable a secure SAN, use the connection table locking feature. In the following example, the connection table is enabled, so no so no other host can login, all management access is disabled and then a single host is enabled to have management access.

```
CLI> SET THIS CONNECTIONS_LOCKED
CLI> SET DISABLE_MANAGERS=ALL
CLI> SET THIS ENABLE_MANAGERS=NEWCON17
```

From this point, the newly enabled host with management privileges or the CLI can be utilized to enable other hosts as management agents.

---

**Note:** The Selective Management Presentation feature only applies to commands received by way of a SCSI SEND\_DIAG command. If the HSG80 receives a SEND\_DIAG command over a disabled management connection, an ILLEGAL\_REQUEST CHECK\_COND will be returned with an ASC=0x91 and ASCQ=0x08. Any command delivered to the HSG80 Serial Port bypasses this constraint and will be processed.

---

## Linking WWIDs for Snap and Clone Units

LUN WWIDs (World Wide Identifiers) for snap and clone units are different each time they are created. This causes more system data records to keep track of the WWIDs as well as script changes at the customer sites. To eliminate this issue, a linked WWID scheme has been created, which keeps the WWIDs of these units constant each time they are created.

The WWID of a LUN is 128bits long, and is composed as follows:

LW3	LW2	LW1	LW0
Controller Node ID	Controller Serial		VSN_Seed

Controller Node ID is the 64-bit Fibre Channel node id of the controller pair.

Controller Serial is the low-order 48 bits of the controller that “initialized” the storage set. The Controller\_Serial is composed from several fields, but the high order 12 bits are a Reserved Field.

VSN\_Seed is a counter that is incremented every time we initialize a storage set.

If the linked WWID is already in use, a unique WWID is allocated, and a message to this effect is displayed. This is not a syntax error and does not cause the command to fail.

## CLI format

CLI> add snapshot\_units <SnapUnit> <UnitName> <SourceName> use\_parent\_wwid

Snapshot Unit - unit number to be assigned to the snap unit.

Unit Name - name of storage set or disk that will become the snap unit.

Source Name - unit number of source storage set or disk.

Example:

CLI> add snap d2 disk10100 d1 use\_parent\_wwid

Example of error message text:

```
CLI> add snap d2 disk10100 d1 use_parent_wwid
A new WWID has been allocated for this unit because the linked
WWID for d2 is already in use.

CLI> run clonew
```

## Implementation Notes

*Add Snap with Linked WWID* - The user has a script that runs every night to create a snapshot, run a backup to tape from the snapshot, then delete the snapshot. Each time this is done, a new WWID is allocated. When the operating system runs out of room for all of these “orphaned” WWIDs, the host system must be rebooted. Therefore the user decides to update the script such that the “add snap” command reads as follows:

```
CLI> add snap d2 disk10100 d1 use_parent_wwid.
```

This results in same WWID being used for the snapshot each night.

*Run Clonew* - Operates the same as “CLI> run clone” with the exception that clonew will use the linked WWID associated with the source unit instead of allocating a new one.

```
CLI> run clonew
```

*Run Clone* - Works the same as v86. In other words a unique WWID is always allocated to the clone unit.

*Clonew of a Snap* - The user wants to clone a snap unit without using any more WWIDs. The clone created from the snap unit will be created using the linked WWID associated with the snap unit.

Exception: A new WWID will be allocated if the snapshot was created using the use\_parent\_wwid switch. Each WWID only has one linked WWID variation. If the linked WWID is already in use, then a new unique WWID is allocated.

*Snap of a Clone* - The user want to snapshot a clone without using any more WWIDs, and issues “CLI> add snap d3 disk10300 d2 use\_parent\_wwid”. D3, the snapshot unit, will be created with the linked WWID associated with d2.

Exception: A new WWID will be allocated if the clone was created using the “run clonew” command. Each WWID only has one linked WWID variation. If the linked WWID is already in use, then a new unique WWID is allocated.

*Manual Clone Creation* - The user has his own set of scripts that create clones, and wants to update them to use linked WWIDs. At some point in the script there will be an “add unit” command. The switch “parent\_wwid=<unit>” must be provided.

For example,

```
CLI> add unit d2 disk10100 parent_wwid=d1
```

would create a unit d2 from device disk10100 whose WWID would be the linked WWID associated with unit d1.

Exception: A new WWID will be allocated if the linked WWID associated with d1 is already in use (for example, a clone or snapshot is already using the linked WWID).

## SMART Error Eject

When a SMART notification is received from a device, it is currently treated as a soft error - the notification is passed to the host and operations continue. A new CLI switch at the controller level changes this behavior. When this switch is enabled, drives in a normalized and redundant set that report a smart error are removed from that set. Smart errors that are reported by drives in a non-redundant or non-normal set continue to handle this condition as a recovered error.

If the smart error eject state is disabled, all smart errors are reported as recovered errors. The recovered error report contains the ASC = 0x5D, the ASC for all smart errors, and the appropriate ASCQ. The default value for this feature is DISABLE.

CLI command syntax:

```
HSG> set this_controller smart_error_eject = [enable|disable]
```

CLI output - feature disabled:

```
AP_TOP> show this
Controller:
    HSG80 ZG02804912 Software V87S-0, Hardware E12
    NODE_ID          = 5000-1FE1-FF00-0090
    ALLOCATION_CLASS = 1
    SCSI_VERSION     = SCSI-3
    Configured for MULTIBUS_FAILOVER with ZG02804288
        In dual-redundant configuration
    Device Port SCSI address 7
    Time: 22-NOV-2001 01:14:32
    Command Console LUN is lun 0 (IDENTIFIER = 99)
    Host Connection Table is NOT locked
    Smart Error Eject Disabled

Host PORT_1:
    Reported PORT_ID = 5000-1FE1-FF00-0093
    PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address         = 011100

Host PORT_2:
    Reported PORT_ID = 5000-1FE1-FF00-0094
    PORT_2_TOPOLOGY = FABRIC (fabric up)
    Address         = 011300
    NOREMOTE_COPY

Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
    CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

Mirrored Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
```

**Battery:**

NOUPS

FULLY CHARGED

Expires: WARNING: UNKNOWN EXPIRATION DATE!

WARNING: AN UNKNOWN NUMBER OF DEEP DISCHARGES HAVE OCCURRED!

CLI Output - feature enabled:

```
AP_TOP> show this
Controller:
    HSG80 ZG02804912 Software V87S-0, Hardware E12
    NODE_ID          = 5000-1FE1-FF00-0090
    ALLOCATION_CLASS = 1
    SCSI_VERSION     = SCSI-3
    Configured for MULTIBUS_FAILOVER with ZG02804288
        In dual-redundant configuration
    Device Port SCSI address 7
    Time: 22-NOV-2001 01:17:47
    Command Console LUN is lun 0 (IDENTIFIER = 99)
    Host Connection Table is NOT locked
    Smart Error Eject Enabled

Host PORT_1:
    Reported PORT_ID = 5000-1FE1-FF00-0093
    PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address         = 011100

Host PORT_2:
    Reported PORT_ID = 5000-1FE1-FF00-0094
    PORT_2_TOPOLOGY = FABRIC (fabric up)
    Address         = 011300
    NOREMOTE_COPY

Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
    CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

Mirrored Cache:
    256 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache

Battery:
```

```
NOUPS
FULLY CHARGED
Expires: WARNING: UNKNOWN EXPIRATION DATE!
WARNING: AN UNKNOWN NUMBER OF DEEP DISCHARGES HAVE OCCURRED!
```

## Error Threshold for Drives

A new limit for driver errors can be set. Once the limit is reached, the drive is removed from any redundant sets to which it belongs and put into the failed set.

Errors counted are medium and recovered errors - there is no need to add hardware errors to this count as the drive fails immediately if a hardware error is encountered.

A set of CLI commands is provided that sets the threshold value and resets the error counters for the drives. This is needed due to the persistent nature of these counters. Since the layered application parsing of the CLI output should not be disturbed, a separate SHOW command is included to see the threshold value.

### CLI Syntax

```
CLEAR_ERRORS DRIVE_ERRORS
```

Enter the DRIVE\_ERRORS command exactly and completely. This command clears the error count, from non-volatile memory, for all devices on the system, setting each to a value of 0. The current value of a devices error count is seen through VTDPY.

```
SET DRIVE_ERROR_THRESHOLD=<value>
```

Enter the DRIVE\_ERROR\_THRESHOLD command exactly and completely. This command sets the threshold error to a value that determines when a drive will be removed and placed in the FAILEDSET. The options for the command are any number ranging from 0 to 999 with a value 0 shutting off the functionality. The user can also input DEFAULT to receive the default error threshold of 700. The entered value is placed into Non-Volatile memory and as such is persistent through failovers and reboots.

```
SHOW DRIVE_ERROR_THRESHOLD
```

This command shows the current value of the drive error threshold for the controller. The output is as follows

```
Drive Error Threshold: <value>
```



# Installing and Configuring the HSG Agent

4

StorageWorks Command Console (SWCC) enables real-time configuration of the storage environment and permits you to monitor and configure the storage connected to the HSG80 controller.

The following information is included in this chapter:

- [Why Use StorageWorks Command Console \(SWCC\)?](#), page 132
- [Installation and Configuration Overview](#), page 133
- [About the Network Connection for the Agent](#), page 134
- [Before Installing the Agent](#), page 136
- [SWCC Features and Components](#), page 132
- [Installing and Configuring the Agent](#), page 137
- [Reconfiguring the Agent](#), page 143
- [Removing the Agent](#), page 145

Refer to Chapter 5 for a description of how to configure a subsystem that uses Fibre Channel fabric topology.

## Why Use StorageWorks Command Console (SWCC)?

StorageWorks Command Console (SWCC) enables you to monitor and configure the storage connected to the HSG80 controller. SWCC consists of Client and Agent.

- The Client provides pager notification and lets you manage your virtual disks. The Client runs on Windows Server 2003 (32-bit), Windows 2000 with Service Pack 3 and 4, and Windows NT 4.0 with Service Pack 6A or above.
- The Agent obtains the status of the storage connected to the controller. It also passes the status of the devices connected to the controller to other computers and provides email notification and error logging.

To receive information about the devices connected to your HSG80 controller over a TCP/IP network, you must install the Agent on a computer that is connected to a controller.

The Agent can also be used as a standalone application without Client. In this mode, which is referred to as Agent only, Agent monitors the status of the subsystem and provides local and remote notification in the event of a failure. A subsystem includes the HSG80 controller and its devices. Remote and local notification can be made by email and/or SNMP messages to an SNMP monitoring program.

**Table 14: SWCC Features and Components**

Features	Agent Required?	Client Required?
Creation of RAID sets: Striped device group (RAID 0) Mirrored device group (RAID 1) Striped mirrored device group (RAID 0+1) Striped parity device group (RAID 3/5) Individual device (JBOD)	Yes	Yes
Monitor multiple subsystems at once	Yes	No
Event logging	Yes	No
Email notification	Yes	No
Pager notification	Yes	Yes

---

**Note:** For serial and SCSI connections, the Agent is not required for creating virtual disks.

---

## Installation and Configuration Overview

**Table 15** provides an overview of the installation.

**Table 15: Installation and Configuration Overview**

Step	Procedure
1.	Verify that your hardware has been set up correctly. See the previous chapters in this guide.
2.	Verify that you have a network connection for the Client and Agent systems. See <a href="#">"About the Network Connection for the Agent" on page 134</a> .
3.	Verify that there is a LUN for communications. This can be either the CCL or a LUN that was created with the CLI. See <a href="#">"What is the Command Console LUN?" on page 38</a> in Chapter 1.
4.	Install the Agent (TCP/IP network connections) on a system connected to the HSG80 controller. See Chapter 3 for agent installation.
5.	Add the name of the Client system to the Agent's list of Client system entries (TCP/IP network connections). This can be done during installation or when reconfiguring the Agent.
6.	Add the name of the Agent system to the Navigation Tree of each Client system that is on the Agent's list of Client system entries (TCP/IP network connections). See Appendix B.
7.	Set up pager notification (TCP/IP network connections). Refer to <a href="#">"Setting Up Pager Notification" in the <i>HP StorageWorks Command Console Version 2.3, User Guide</i></a> .

## About the Network Connection for the Agent

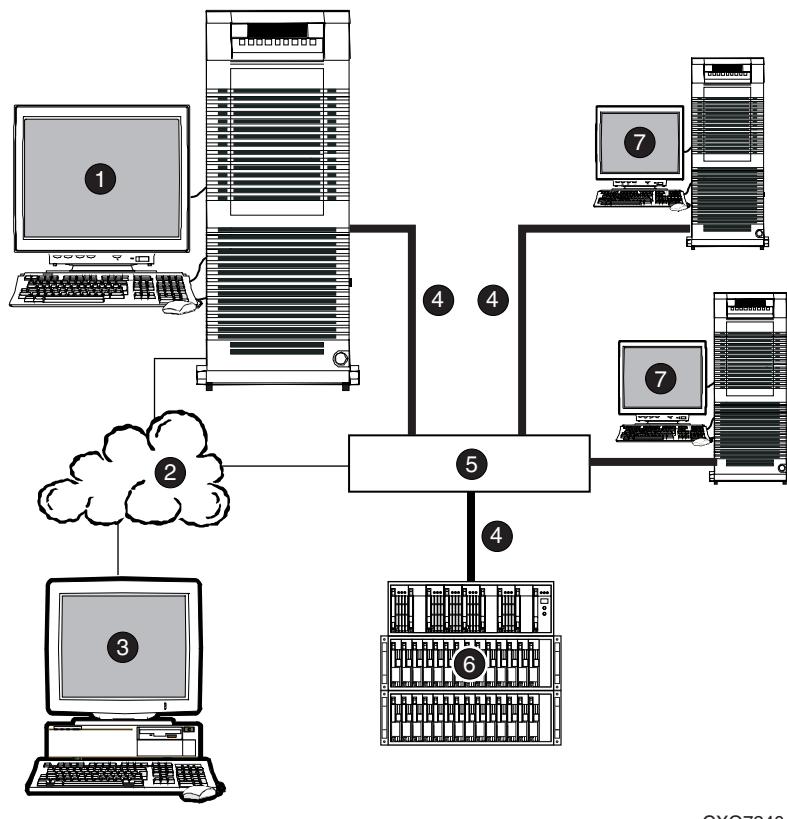
The network connection, shown in [Figure 35](#), displays the subsystem connected to a hub or a switch. SWCC can consist of any number of Clients and Agents in a network. However, it is suggested that you install only one Agent on a computer.

By using a network connection, you can configure and monitor the subsystem from anywhere on the LAN. If you have a WAN or a connection to the Internet, monitor the subsystem with TCP/IP.

---

**Note:** SWCC does not support the dynamic host configuration protocol (DHCP) or the Windows Internet Name Service (WINS).

---



CX07240A

① Agent system (has the Agent software)	⑤ Hub or switch
② TCP/IP network	⑥ HSG80 controller and its device subsystem
③ Client system (has the Client software)	⑦ Servers
④ Fibre Channel cable	

**Figure 35: An example of a network connection**

## Before Installing the Agent

The Agent requires the minimum system requirements, as defined in the release notes for your operating system. The program is designed to operate with the Client version 2.5 on Windows 2000, Windows NT, or Windows Server 2003 (32-bit). Verify that your system meets the minimum requirements by completing the following steps.

1. Login as root (superuser). Agent installations on Tru64 UNIX must be done locally. Do not attempt to install the Agent over the network.
2. Remove previous versions of the Agent from your computer.
3. Read the release notes.

---

**Note:** Refer to Appendix C for further detail on how to setup the SWCC Agent in a TruCluster Environment.

---



**Caution:** Do not install the Agent by using Dataless Management Services (DMS) or Remote Installation Services (RIS) server installations.

---

## Downloading the Host Kit Software From the Web

The host kit software is available for download. You can save the software to your computer or create a CD-ROM. Platform kit software is stored on the download web site based on operating system. Follow the steps below to obtain the software from the web site.

1. Go to  
<http://h18006.www1.hp.com/products/storageworks/ma8kema12k/kits.html>.
2. Select the kit for download.
3. Select the folder in which you want to save the file when presented with the Save As dialog box.

---

**Note:** The folder you select must be large enough to store the OS kit selected. File sizes range from 862KB to 51 MB depending on the operating system selected.

---

The platform kit software is available as an .iso or .zip file. After you download it to your computer, your options are:

- Create a CD-ROM using the .iso file.
- Copy the .zip file to the server that runs the platform kit software. Uncompress the file and run the install program on the server.

## Installing and Configuring the Agent

The following instructions assume that you have a directory /mnt to which you can mount the CD-ROM. If you do not, you must create a mount point and replace /mnt in the following sequence with the mount point you created. It also assumes your CD-ROM device is /dev/rz4c. If not, replace /dev/rz4c with the actual CD-ROM device.

1. Obtain the following information:
  - Names of systems you want the Agent to contact
  - Names of subsystems connected to the host
  - Email addresses for the Agent to send status updates
2. Insert the CD-ROM into a computer.
3. Enter one of the following commands at the command prompt:

— For V4.0x enter:

```
# mount -r -t cdfs -o rrip /dev/rz6c /mnt
```

(Substituting rz6c if necessary for your CD-ROM)

— For V5.1x enter:

```
# mount -r -t cdfs /dev/disk/cdrom0c /mnt
```

(Substituting cdrom0c if necessary for your CD-ROM)

4. Change directories on the CD-ROM by entering:

```
# cd /mnt/agent
```



**Caution:** The version of *Compaq Tru64 UNIX* that you are using, either V4.0x or V5.1x will determine which *Device Special File Name* format you will need to enter.

---

5. To run the installation program, enter the following at the command prompt:

```
# setld -l .
```

---

**Note:** The -l is a lowercase L, and the period must appear at the end.

---

You are asked if you want to install the listed subsets.

6. Enter the corresponding number of the software and then press the **Enter** key.

## Configure Client System Options

7. Select option y and then press the **Enter** key.

As the installation continues, you see several system messages, telling you that the subset has been installed and is being loaded. The computer asks you to add a client system entry.

8. Enter the name of the client system that you want to receive updates from this Agent. Press the **Enter** key.

---

**Note:** Enter your most important client system first. Enter client systems that are connected infrequently last. The Agent puts the client systems entered first at the top of its list. It first contacts the client system located at the top of the list.

---

9. From the access options menu, enter an access level for the client system.

The access privilege level controls the client system's level of access to the subsystems. The following explains each access option:

**Table 16: Client System Access Options**

Options	SWCC Function
0 = Overall Status	Can use the Client to add a system to a Navigation Tree, set up a pager, and view properties of the controller and the system Cannot use Client to open a Storage Window
1 = Detailed Status	Can use the Client to open a Storage Window, but you cannot make modifications in that window
2 = Configuration and Status	Can use the Client to make changes in a Storage Window to modify a subsystem configuration

10. Press **Enter**.

A menu for selecting a client system notification scheme appears:

11. From the menu, enter a notification scheme for the client system: **0, 1, 2, or 3**.

The notification scheme defines the network protocol to be used by the Agent when notifying the selected client system of a change in the state of a subsystem. The following describes how the Transmission Control Protocol/Internet Protocol (TCP/IP) and the Simple Network Management Protocol (SNMP) work with SWCC.

**Table 17: Client System Notification Options**

Options	SWCC Function
Transmission control Protocol/Internet Protocol (TCP/IP)	Automatically updates the Storage Window of subsystem changes provided AES is running Required for Windows NT event logging and pager notification If you do not select TCP/IP, you will need to refresh the Storage Window (depending on the subsystem) to obtain the latest status of a subsystem.
Simple Network Management Protocol (SNMP)	Requires you to use an SNMP-monitoring program to view SNMP traps

12. Press **Enter**.

The computer asks you if the entered information is correct.

13. Select option **y** and then press the **Enter** key.

A message, asking if you would like to add another Client, appears.

14. To stop adding client systems, select option **n** and press the **Enter** key.

You are asked for a password, which is required to do configurations within the Client software. If an old password is found, you are asked if you want to use it.

## Enter Subsystem Information

15. Enter your case-sensitive password that has 4 to 16 characters, and press the **Enter** key. You are asked to retype the password.

16. Retype the password and press the **Enter** key.

Once the password has been entered, the system scans for subsystems. You will be asked for the name and monitoring interval of each RAID subsystem found.

17. Enter the subsystem name in lowercase letters, and press the **Enter** key. You are asked for the monitoring interval. The monitoring interval is the rate at which the Agent queries the specified subsystem for status.

18. Enter the monitoring interval in seconds, and press the **Enter** key. The monitoring interval determines how often the Agent polls the subsystems. You are asked if the displayed information is correct.

19. If the displayed information is correct, select option **y** and press the **Enter** key.

You are asked if you want Email enabled.

## Configure Email Notification Details

20. To enable the Email, select option **y** and press the **Enter** key. The software asks for the address of the person to notify.

21. Enter the address of the person to notify, and press the **Enter** key.

22. From the displayed menu, enter the notification level: **1**, **2**, or **3**.

[Table 18](#) provides the definitions of the Email notification options:

**Table 18: Definitions of Email Notification Options**

Term	Definition
Critical Errors	Notifies you of errors that would prevent you from accessing data; for example, a disk failing or a client system not available on the network.
Warnings	Notifies you when something is broken, but its breakage does not prevent you from accessing data; for example, the RAID degrading or losing a fan.
Information	Provides messages, but they do not indicate that something is broken. Examples of informational messages are the following: an Agent startup message or a message saying that an error has been resolved.

23. Press **Enter**.

The software asks if the displayed information is correct.

24. If the displayed information is correct, select option **y** and press the **Enter** key.

A message, asking if you would like to add another person to notify, appears.

25. Select option **n** to end the dialog or **y** to add another person to notify. Press **Enter**.

After you select option **n**, the system will start the new Agent. If the installation script detects a problem at Agent startup, it will alert you.

You are asked if you are running Agent in a TruCluster environment.

---

**Note:** SWCC in A TruCluster environment (Appendix C) goes into detail on how to setup the SWCC Agent in a TruCluster Environment.

---

## Final System Configuration

26. Do one of the following:

- If you are running the Agent in cluster, select option **y**. The software is configured to scan the SCSI-bus for subsystems at startup.
- If you are not running the Agent in a cluster, select option **n**. The software is configured to look at the **storage.ini** file at startup time. The **storage.ini** file contains the listing of your subsystems.

---

**Note:** StorageWorks recommends that if you are running Tru64 UNIX 5.1x, especially within a fabric, configure the software to scan the SCSI-bus at startup. When asked if running the Agent in a cluster, answer **y**.

---

The system starts the new Agent.

27. Enter the following and press the **Enter** key to unmount the CD-ROM:

```
# cd /
# umount /mnt
```

## Reconfiguring the Agent

You can change your configuration using the SWCC Agent Configuration menu. To access this menu, enter the following command:

```
# /usr/opt/SWCC520/scripts/swcc_config
```

The following is an example of the menu::

```
SWCC Agent Configuration Utility
-----
-----
Options Available Are:
1) Add/Delete Client PC
Information.
2) Modify Storage Subsystem
Information.

3) Add/Delete Storage Subsystem
Information
4) Change SWCC Agent Password.
5) Turn Email Notification ON/OFF.
6) Add/Delete Email Notification
Users.
7) Restart Agent With Changes.
8) Enable/Disable Agent.

9) Scan/Don't Scan on Startup
10) Cluster/No Cluster
Enter a number or 'q' to QUIT:
```



**Caution:** After you make a change to the configuration, such as adding a client system, you must stop and then start the Agent for your changes take effect. When you stop and then start the Agent, the Storage Windows for the subsystems connected to the agent system lose their connection. To regain that connection, close and then reopen the Storage Windows connected to the agent system after you restart the Agent.

---

**Table 19: Information Needed to Configure Agent**

Term/Procedure	Description
Adding a Client system entry	For a client system to receive updates from the Agent, you must add it to the Agent's list of client system entries. The Agent will only send information to client system entries that are on this list. In addition, adding a client system entry allows you to access the Agent system from the Navigation Tree on that Client system.
Adding a subsystem entry	You need to tell the Agent the subsystem that it needs to monitor.
Client system	Network names for the computers on which the Client software runs.
Client system access options	<p>The access privilege level controls the Client system's level of access to the subsystems.</p> <p>0=No Access-Can use the Client software to add a system to a Navigation Tree, set up a pager, and view properties of the controller and the system. You cannot use Client to open a Storage Window.</p> <p>1=Show Level Access-Can use the Client software to open a Storage Window, but you cannot make modifications in that window.</p> <p>2=Storage Subsystem Configuration Capability-Can use the Client software to make changes in a Storage Window to modify a subsystem configuration.</p>
<p>Client system notification options</p> <p><b>Note:</b> For all of the client system notification options, local notification is available through an entry in the system error log file and Email (provided that Email notification in PAGEMAIL.COM has not been disabled).</p>	<p>0 = No Error Notification-No error notification is provided over network.</p> <p>1 = Notification via a TCP/IP Socket (Transmission Control Protocol/Internet Protocol)-Updates the Storage Window of subsystem changes provided AES is running. Required for Windows NT event logging and pager notification. If you do not select TCP/IP, you will need to refresh the Storage Window to obtain the latest status of a subsystem.</p> <p>2 = Notification via the SNMP protocol (Simple Network Management Protocol)-Requires you to use an SNMP-monitoring program to view SNMP traps.</p> <p>3 = Notification via both TCP/IP and SNMP-Combination of options 1 and 2.</p>

**Table 19: Information Needed to Configure Agent (Continued)**

Term/Procedure	Description
Deleting a client system entry	When you remove a client system from the Agent's list, you are instructing the Agent to stop sending updates to that client system. In addition, you will be unable to access this agent system from the Navigation Tree.
Email notification	Modify file pagemail.com in directory <b>sys\$sysdevice:[swcc\$agent]</b> . When an error is logged, the Agent executes the PAGEMAIL.COM command. You can modify this file for Agent to log errors in a log file and/or change the account to which Agent sends messages. You can also modify for which level of errors you will be notified. Client does not need to be running to perform these actions.
Monitoring interval in seconds	How often the subsystem is monitored.
Password	It must be a text string that has 4 to 16 characters. It can be entered from the client system to gain configuration access. Accessing the SWCC Agent Configuration menu can change it.

## Removing the Agent



**Caution:** Do not uninstall the Agent if you want to preserve configuration information. If you only want to install an upgrade, stop the Agent, and then install the new version. Older versions will be automatically removed before the update, but all configuration information will be preserved.

1. Enter the following and press the **Enter** key:

```
# setld -d SWCC520
```

The procedure to remove the Agent asks the following:

Are you sure you want to delete all Agent data?

2. Do one of the following:

- Select option **y** and then press the **Enter** key. All configuration files for the Agent is deleted.

- Select option **n** and press the **Enter** key. The computer saves the configuration data, which includes the following:
  - Subsystem list of entries
  - Client system access options
  - Password
  - Email notification

You are able to use the data if you reinstall the Agent or install a newer version of the Agent.

# FC Configuration Procedures

5

This chapter describes procedures to configure a subsystem that uses Fibre Channel (FC) fabric topology. In fabric topology, the controller connects to its hosts through switches.

The following information is included in this chapter:

- [Establishing a Local Connection](#), page 148
- [Setting Up a Single Controller](#), page 149
- [Setting Up a Controller Pair](#), page 156
- [Configuring Devices](#), page 163
- [Configuring Storage Containers](#), page 163
- [Assigning Unit Numbers and Unit Qualifiers](#), page 169
- [Configuration Options](#), page 170
- [Verifying Storage Configuration from Host](#), page 175

Use the command line interpreter (CLI) or *StorageWorks* Command Console (SWCC) to configure the subsystem. This chapter uses CLI to connect to the controller. To use SWCC for configuration, see the SWCC online help for assistance.

---

**Note:** These configuration procedures assume that controllers and cache modules are installed in a fully functional and populated enclosure and that the PCMCIA cards are installed.

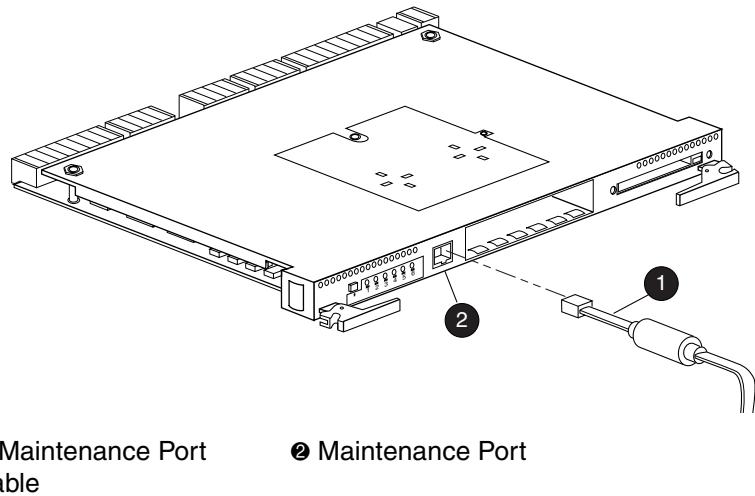
---

To install a controller or cache module and the PCMCIA card, see the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Maintenance and Service Guide*.

## Establishing a Local Connection

A local connection is required to configure the controller until a command console LUN (CCL) is established using the CLI. Communication with the controller can be through the CLI or SWCC.

The maintenance port, shown in [Figure 36](#), provides a way to connect a maintenance terminal. The maintenance port can be an EIA-423 compatible terminal or a computer running a terminal emulator program. The maintenance port accepts a standard RS-232 jack. The maintenance port cable shown in [Figure 36](#) has a 9-pin connector molded onto the end for a PC connection. If you need a terminal connection or a 25-pin connection, you can order optional cabling.



**Figure 36: Maintenance port connection**



**Caution:** The maintenance port generates, uses, and can radiate radio-frequency energy through its cables. This energy may interfere with radio and television reception. Disconnect all maintenance port cables when not communicating with the controller through the local connection.

## Setting Up a Single Controller

### Powering On and Establishing Communication

1. Connect the computer or terminal to the controller, as shown in [Figure 36](#). The connection to the computer is through the COM1 or COM2 port.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Verify that the computer or terminal is configured as follows:
  - 9600 baud
  - 8 data bits
  - 1 stop bit
  - no parity
  - no flow control
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

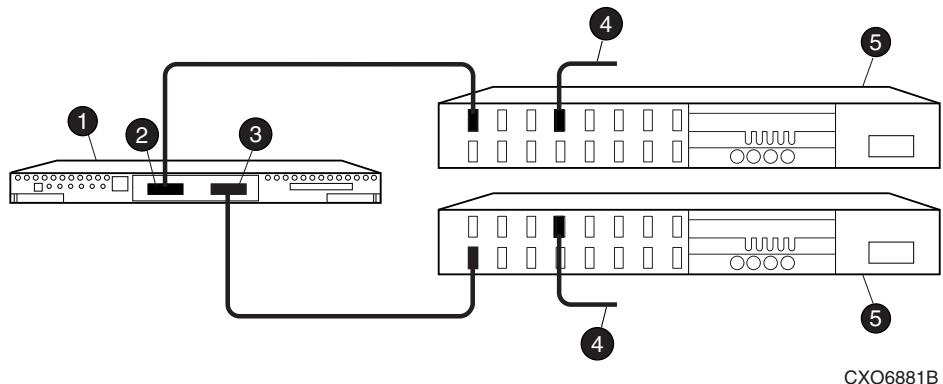
### Cabling a Single Controller

The cabling for a single controller is shown in [Figure 37](#).

---

**Note:** It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure ("Configuring a Single Controller Using CLI" on page 150).

---



- ① Controller
- ② Host port 1
- ③ Host port 2
- ④ Cable from the switch to the host Fibre Channel adapter
- ⑤ FC switch

**Figure 37: Single controller cabling**

## Configuring a Single Controller Using CLI

To configure a single controller using CLI involves the following processes:

- Verifying the Node ID and Checking for Any Previous Connections.
- Configuring Controller Settings.
- Restart the Controller.
- Setting Time and Verifying all Commands.
- Plugging in the FC Cable and Verify Connections.
- Repeating Procedure for Each Host Adapter.
- Verifying Installation.

### Verifying the Node ID and Checking for Any Previous Connections

1. Enter a `SHOW THIS` command to verify the node ID:

```
SHOW THIS
```

See “[Worldwide Names \(Node IDs and Port IDs\)](#)” on page 55 for the location of the sticker.

The node ID is located in the third line of the SHOW THIS result:

```
HSG80> SHOW THIS
Controller:
HSG80 ZG80900583 Software V8.8, Hardware E11
NODE_ID          = 5000-1FE1-0001-3F00
ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 5.

If the node ID is all zeroes, enter node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

```
SET THIS NODE_ID=NNNN-NNNN-NNNN-NNNN nn
```

Where:NNNN-NNNN-NNNN-NNNN is the node ID, and *nn* is the checksum.

2. When using a controller that is not new from the factory, enter the following command to take it out of any failover mode that may have been configured previously:

```
SET NOFAILOVER
```

If the controller did have a failover mode previously set, the CLI may report an error. Clear the error with this command:

```
CLEAR_ERRORS CLI
```

3. Enter the following command to remove any previously configured connections:

```
SHOW CONNECTIONS
```

A list of named connections, if any, is displayed.

4. Delete these connections by entering the following command:

```
DELETE !NEWCON01
```

Repeat the `Delete` command for each of the listed connections. When completed, no connections will be displayed.

## Configuring Controller Settings

5. Set the SCSI version using the following command syntax:

```
SET THIS SCSI_VERSION=SCSI-2, or
```

```
SET THIS SCSI_VERSION=SCSI-3
```

**Note:** If SCSI-2 is selected, you must disable CCL using the command:

```
SET THIS NOCOMMAND_CONSOLE_LUN
```

6. Tru64 UNIX V5.1x can use both SCSI-2 or SCSI-3. Assign an identifier for the communication LUN (also called the command console LUN, or CCL). The CCL must have a unique identifier that is a decimal number in the range 1 to 32767, and which is different from the identifiers of all units. Use the following syntax:

```
SET THIS IDENTIFIER=N
```

Identifier must be unique among all the controllers attached to the fabric within the specified allocation class.

7. Set the topology for the controller. If both ports are used, set topology for both ports:

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

If the controller is not factory-new, it may have another topology set, in which case these commands will result in an error message. If this happens, take both ports offline first, then reset the topology:

```
SET THIS PORT_1_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_2_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

## Restarting the Controller

8. Restart the controller, using the following command:

```
RESTART THIS
```

## Setting Time and Verifying all Commands

1. Set the time on the controller by entering the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

2. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer “**Y**”:

```
Do you intend to replace this controller's cache battery? Y/N  
[N] Y
```

FRUTIL will print out a procedure, but will not give you a prompt. Ignore the procedure and press the **Enter** key.

3. Set up any additional optional controller settings, such as changing the CLI prompt. See the *SET THIS CONTROLLER/OTHER CONTROLLER* command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for the format of optional settings.
4. Verify that all commands have taken effect. Use the following command:

```
SHOW THIS
```

Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology.

The following sample is a result of a **SHOW THIS** command, with the areas of interest in **bold**.

Controller:

```
HSG80 ZG94214134 Software V8.8, Hardware 0000
NODE_ID          = 5000-1FE1-0007-9750
ALLOCATION_CLASS = 0
SCSI_VERSION     = SCSI-3
Configured for dual-redundancy with ZG9421461
In dual-redundant configuration
```

Device Port SCSI address 7

```
Time: 10-Mar-2002:12:30:34
Command Console LUN is disabled
Smart Error Eject Disabled
```

Host PORT\_1:

```
Reported PORT_ID = 5000-1FE1-0007-9751
PORT_1_TOPOLOGY = FABRIC (fabric up)
Address         = 7D4000
```

Host PORT\_2:

```
Reported PORT_ID = 5000-1FE1-0007-9752
PORT_2_TOPOLOGY = FABRIC (standby)
Address         = 210513
```

NOREMOTE\_COPY

Cache:

```
512 megabyte write cache, version 0022
Cache is GOOD
No unflushed data in cache
CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

Mirrored Cache:

Not enabled

Battery:

```
NOUPS
FULLY CHARGED
Expires: 25-JUN-2003
.....
```

5. Turn on the switches, if not done previously.

If you want to communicate with the Fibre Channel switches through Telnet, set an IP address for each switch. See the manuals that came with the switches for details.

## Plugging in the FC Cable and Verify Connections

6. Plug the Fibre Channel cable from the first host bus adapter into the switch. Enter the SHOW CONNECTIONS command to view the connection table:

```
SHOW CONNECTIONS
```

7. Rename the connections to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

For a recommended naming convention, see “[Naming Connections](#)” on page 39.

8. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=TRU64_UNIX
```

9. Verify the changes:

```
SHOW CONNECTIONS
```

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

## Repeating Procedure for Each Host Adapter

10. Repeat step 7, 8, and 9 for each of that adapter’s host connections, or delete the unused connections from the table.

11. For each host adapter, repeat steps 6 through 10.

## Verifying Installation

To verify installation for your Tru64\_UNIX host, enter the following commands:

### For V5.1x Use

```
file /dev/cport/ *
```

Your host computer should report one CCL device special file for each HSG80 configured.

## Setting Up a Controller Pair

### Powering Up and Establishing Communication

1. Connect the computer or terminal to the controller as shown in [Figure 36](#). The connection to the computer is through the COM1 or COM2 ports.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
  - 9600 baud
  - 8 data bits
  - 1 stop bit
  - no parity
  - no flow control
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

### Cabling a Controller Pair

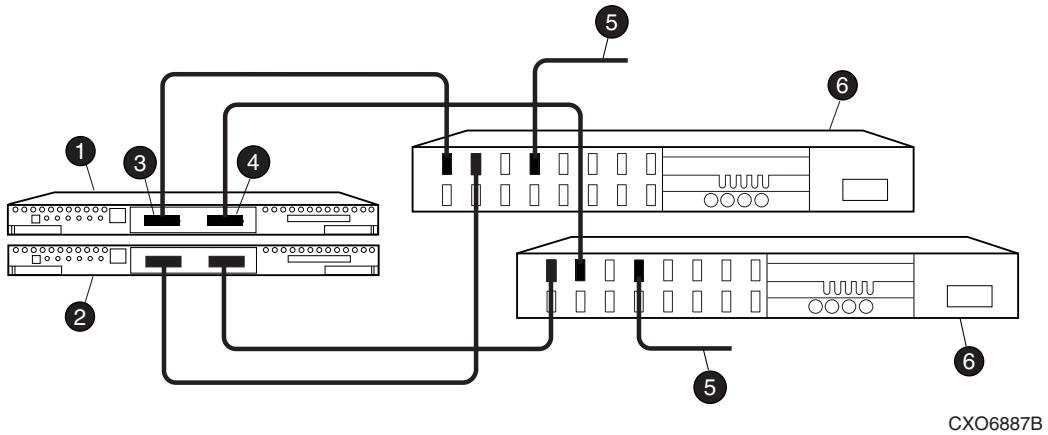
The cabling for a controller pair is shown in [Figure 38](#).

---

**Note:** It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure ("Configuring a Controller Pair Using CLI" on page 157).

---

Figure 38 shows a controller pair with failover cabling showing one HBA per server with HSG80 controller in transparent failover mode.



❶ Controller A	❷ Controller B	❸ Host port 1	❹ Host port 2	❺ Cable from the switch to the host FC adapter	❻ FC switch
----------------	----------------	---------------	---------------	--	-------------

Figure 38: Controller pair failover cabling

## Configuring a Controller Pair Using CLI

To configure a controller pair using CLI:

- Configuring Controller Settings.
- Restarting the Controller.
- Setting Time and Verifying All Commands.
- Plugging in the FC Cable and Verify Connections.
- Repeating Procedure for Each Host Adapter.
- Verifying Installation.

Enter a `SHOW THIS` command to verify the node ID:`SHOW THIS`

See “[Worldwide Names \(Node IDs and Port IDs\)](#)” on page 55, for the location of the sticker.

The node ID is located in the third line of the SHOW THIS result:

```
HSG80> show this
Controller:
HSG80 ZG80900583 Software V8.8, Hardware E11
NODE_ID          = 5000-1FE1-0001-3F00
ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 5.

If the node ID is all zeroes, enter the node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

```
SET THIS NODE_ID=NNNN-NNNN-NNNN-NNNN nn
```

Where: *NNNN-NNNN-NNNN-NNNN* is the node ID and *nn* is the checksum.

6. If the controller is not new from the factory, enter the following command to take it out of any failover mode that may have been previously configured:

```
SET NOFAILOVER
```

If the controller did have a failover mode previously set, the CLI may report an error. Clear the error with this command:

```
CLEAR_ERRORS CLI
```

7. Enter the following command to remove any previously configured connections:

```
SHOW CONNECTIONS
```

A list of named connections, if any, is displayed.

8. Delete these connections by entering the following command:

```
DELETE !NEWCON01
```

Repeat the Delete command for each of the listed connections. When completed, no connections will be displayed.

## Configuring Controller Settings

9. Set the SCSI version to SCSI-3 for Tru64 UNIX V5.1x only, using the following command:

```
SET THIS SCSI_VERSION=SCSI-3
```

---

**Note:** Setting the SCSI version to SCSI-3 does not make the controller fully compliant with the SCSI-3 standards.

---

10. Set the topology for the controller. If both ports are used, set topology for both ports:

```
SET THIS PORT_1_TOPOLOGY=FABRIC  
SET THIS PORT_2_TOPOLOGY=FABRIC
```

If the controller is not factory-new, it may have another topology set, in which case these commands will result in an error message. If this happens, first take both ports offline, then reset the topology:

```
SET THIS PORT_1_TOPOLOGY=OFFLINE  
SET THIS PORT_2_TOPOLOGY=OFFLINE  
SET THIS PORT_1_TOPOLOGY=FABRIC  
SET THIS PORT_2_TOPOLOGY=FABRIC
```

## Restarting the Controller

11. Restart the controller, using the following command:

```
RESTART THIS
```

It takes about a minute for the CLI prompt to come back after a RESTART command.

## Setting Time and Verifying All Commands

12. Set the time on the controller by entering the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

13. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer “Y”:

```
Do you intend to replace this controller's cache battery? Y/N  
[N] Y
```

FRUTIL will print out a procedure, but will not give you a prompt. Ignore the procedure and press **Enter**.

14. Set up any additional optional controller settings, such as changing the CLI prompt. See the *SET THIS CONTROLLER/OTHER CONTROLLER* command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for the format of optional settings. Perform this step on both controllers.
15. Verify that all commands have taken effect by entering the following command:

SHOW THIS

16. Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology. The following display is a sample result of a **SHOW THIS** command, with the areas of interest in bold.

Controller:

```
HSG80 ZG94214134 Software V8.8, Hardware 0000
  NODE_ID          = 5000-1FE1-0007-9750
  ALLOCATION_CLASS = 0
  SCSI_VERSION     = SCSI-3
  Configured for dual-redundancy with ZG9421461
  In dual-redundant configuration

  Device Port SCSI address 7
    Time: 10-Mar-2002:12:30:34
    Command Console LUN is disabled
    Smart Error Eject Disabled

  Host PORT_1:
    Reported PORT_ID = 5000-1FE1-0007-9751
    PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address         = 7D4000

  Host PORT_2:
    Reported PORT_ID = 5000-1FE1-0007-9752
    PORT_2_TOPOLOGY = FABRIC (standby)
    Address         = 210513

  NOREMOTE_COPY

  Cache:
    512 megabyte write cache, version 0022
    Cache is GOOD
    No unflushed data in cache
    CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

  Mirrored Cache:
    Not enabled

  Battery:
    NOUPS
    FULLY CHARGED
    Expires: 25-JUN-2003
```

17. Turn on the switches if not done previously.

If you want to communicate with the FC switches through Telnet, set an IP address for each switch. See the manuals that came with the switches for details.

## Plugging in the FC Cable and Verify Connections

18. Plug the FC cable from the first host adapter into the switch. Enter a `SHOW CONNECTIONS` command to view the connection table:

```
SHOW CONNECTIONS
```

The first connection will have one or more entries in the connection table. Each connection will have a default name of the form `!NEWCONxx`, where `xx` is a number representing the order in which the connection was added to the connection table.

For a description of why plugging in one adapter can result in multiple connections, see “[Numbers of Connections](#)” on page 40.

19. Rename the connections to something meaningful to the system and easy to remember. For example, to assign the name `ANGEL1A1` to connection `!NEWCON01`, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

HP recommends using a naming convention, see “[Naming Connections](#)” on page 39.

20. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=TRU64_UNIXAIX_CAMBEX
```

21. Verify the changes:

```
SHOW CONNECTIONS
```

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

## Repeating Procedure for Each Host Adapter Connection

22. Repeat steps 19, 20, and 21 for each of that adapter’s host connections or delete the unwanted connections from the table.
23. For each host adapter, repeat steps 18 through 22.

## Verifying Installation

To verify installation for your Tru64\_UNIX host, enter one of the following commands:

## For V5.1x Use

```
file /dev/cport/ *
```

Your host computer should report one CCL device special file for each HSG80 configured.

## Configuring Devices

The disks on the device bus of the HSG80 can be configured manually or with the CONFIG utility. The CONFIG utility is easier. Invoke CONFIG with the following command:

```
RUN CONFIG
```



**WARNING:** It is highly recommended to use the CONFIG utility only at reduced I/O loads.

---

CONFIG takes about two minutes to discover and to map the configuration of a completely populated storage system.

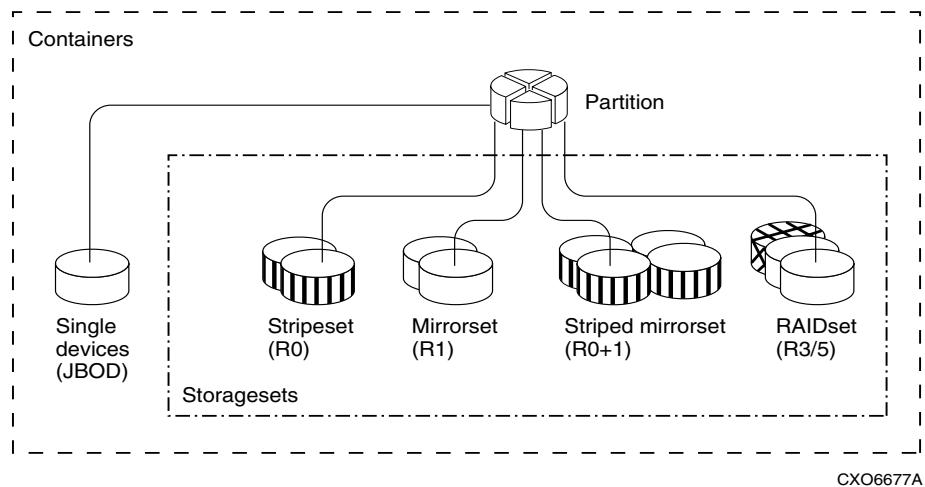
## Configuring Storage Containers

For a technology refresher on this subject, refer to “[Choosing a Container Type](#)” on page 72.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in [Figure 39](#). The independent disks and the selected storageset may also be partitioned.

The following configurations are detailed in the following sections:

- [Configuring a Stripeset](#), page 164
- [Configuring a Mirrorset](#), page 165
- [Configuring a RAIDset](#), page 165
- [Configuring a Striped Mirrorset](#), page 166
- [Configuring a Single-Disk Unit \(JBOD\)](#), page 167
- [Configuring a Partition](#), page 167



**Figure 39: Storage container types**

## Configuring a Stripeset

1. Create the stripeset by adding its name to the controller's list of storagesets and by specifying the disk drives it contains. Use the following syntax:

```
ADD STRIPESET STRIPESET-NAME DISKNNNNN DISKNNNNN.....
```

2. Initialize the stripeset, specifying any desired switches:

```
INITIALIZE STRIPESET-NAME SWITCHES
```

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

3. Verify the stripeset configuration:

```
SHOW STRIPESET-NAME
```

4. Assign the stripeset a unit number to make it accessible by the hosts. See “[Assigning Unit Numbers and Unit Qualifiers](#)” on page 169.

For example:

The commands to create Stripe1, a stripeset consisting of three disks (DISK10000, DISK20000, and DISK10100) and having a chunksize of 128:

```
ADD STRIPESET STRIPE1 DISK10000 DISK20000 DISK30000
```

```
INITIALIZE STRIPE1 CHUNKSIZE=128
```

```
SHOW STRIPE1
```

## Configuring a Mirrorset

1. Create the mirrorset by adding its name to the controller's list of storage sets and by specifying the disk drives it contains. Optionally, you can append mirrorset switch values:

```
ADD MIRRORSET MIRRORSET-NAME DISKNNNNN DISKNNNNN SWITCHES
```

---

**Note:** See the ADD MIRRORSET command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for a description of the mirrorset switches.

---

2. Initialize the mirrorset, specifying any desired switches:

```
INITIALIZE MIRRORSET-NAME SWITCHES
```

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

3. Verify the mirrorset configuration:

```
SHOW MIRRORSET-NAME
```

4. Assign the mirrorset a unit number to make it accessible by the hosts. See “[Assigning Unit Numbers and Unit Qualifiers](#)” on page 169.

For example:

The commands to create Mirr1, a mirrorset with two members (DISK10000 and DISK20000), and to initialize it using default switch settings:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000
```

```
INITIALIZE MIRR1
```

```
SHOW MIRR1
```

## Configuring a RAIDset

1. Create the RAIDset by adding its name to the controller's list of storage sets and by specifying the disk drives it contains. Optionally, you can specify RAIDset switch values:

```
ADD RAIDSET RAIDSET-NAME DISKNNNNN DISKNNNNN DISKNNNNN SWITCHES
```

**Note:** See the ADD RAIDSET command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for a description of the RAIDset switches.

---

2. Initialize the RAIDset, specifying any desired switches:

```
INITIALIZE RAIDSET-NAME SWITCH
```

---

**Note:** HP recommends that you allow initial reconstruct to complete before allowing I/O to the RAIDset. Not doing so may generate forced errors at the host level. To determine whether initial reconstruct has completed, enter SHOW RAIDSET FULL.

---

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

3. Verify the RAIDset configuration:

```
SHOW RAIDSET-NAME
```

4. Assign the RAIDset a unit number to make it accessible by the hosts. See “[Assigning Unit Numbers and Unit Qualifiers](#)” on page 169.

For example:

The commands to create RAID1, a RAIDset with three members (DISK10000, DISK20000, and DISK10100) and to initialize it with default values:

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK30000
INITIALIZE RAID1
SHOW RAID1
```

## Configuring a Striped Mirrorset

1. Create, but do not initialize, at least two mirrorsets.

See “[Configuring a Mirrorset](#)” on page 165.

2. Create a stripeset and specify the mirrorsets it contains:

```
ADD STRIPESET STRIPESET-NAME MIRRORSET-1
MIRRORSET-2....MIRRORSET-N
```

3. Initialize the striped mirrorset, specifying any desired switches:

```
INITIALIZE STRIPESET-NAME SWITCH
```

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

4. Verify the striped mirrorset configuration:

```
SHOW STRIPESET-NAME
```

5. Assign the stripeset mirrorset a unit number to make it accessible by the hosts. See “[Assigning Unit Numbers and Unit Qualifiers](#)” on page 169.

For example:

The commands to create Stripe1, a striped mirrorset that comprises Mirr1, Mirr2, and Mirr3, each of which is a two-member mirrorset:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000
ADD MIRRORSET MIRR2 DISK20100 DISK10100
ADD MIRRORSET MIRR3 DISK10200 DISK20200
ADD STRIPESET STRIPE1 MIRR1 MIRR2 MIRR3
INITIALIZE STRIPE1
SHOW STRIPE1
```

## Configuring a Single-Disk Unit (JBOD)

1. Initialize the disk drive, specifying any desired switches:

```
INITIALIZE DISK-NAME SWITCHES
```

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

2. Verify the configuration by entering the following command:

```
SHOW DISK-NAME
```

3. Assign the disk a unit number to make it accessible by the hosts. See “[Assigning Unit Numbers and Unit Qualifiers](#)” on page 169.

## Configuring a Partition

1. Initialize the storage set or disk drive, specifying any desired switches:

```
INITIALIZE STORAGESET-NAME SWITCHES
```

or

```
INITIALIZE DISK-NAME SWITCHES
```

See “[Specifying Initialization Switches](#)” on page 89 for a description of the initialization switches.

2. Create each partition in the storageset or disk drive by indicating the partition's size. Also specify any desired switch settings:

```
CREATE_PARTITION STORAGESSET-NAME SIZE=N SWITCHES
```

or

```
CREATE_PARTITION DISK-NAME SIZE=N SWITCHES
```

where N is the percentage of the disk drive or storageset that will be assigned to the partition. Enter SIZE=LARGEST, on the last partition only, to let the controller assign the largest free space available to the partition.

---

**Note:** See the *CREATE\_PARTITION* command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide* for a description of the partition switches.

---

3. Verify the partitions:

```
SHOW STORAGESSET-NAME
```

or

```
SHOW DISK-NAME
```

The partition number appears in the first column, followed by the size and starting block of each partition.

4. Assign the partition a unit number to make it accessible by the hosts. See ["Assigning Unit Numbers and Unit Qualifiers"](#) on page 169.

For example:

The commands to create RAID1, a three-member RAIDset, then partition it into two storage units are shown below.

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK30000
INITIALIZE RAID1
CREATE_PARTITION RAID1 SIZE=25
CREATE_PARTITION RAID1 SIZE=LARGEST
SHOW RAID1
```

## Assigning Unit Numbers and Unit Qualifiers

Each storageset, partition, or single (JBOD) disk must be assigned a unit number for the host to access. As the units are added, their properties can be specified through the use of command qualifiers, which are discussed in detail under the ADD UNIT command in the *HP StorageWorks HSG60 and HSG80 Array Controller and Array Controller Software Command Line Interface (CLI) Reference Guide*.

Because of different SCSI versions, refer to the section “[Assigning Unit Numbers Depending on SCSI\\_VERSION](#)” on page 47. The choice for SCSI\_VERSION effects how certain unit numbers and host connection offsets interact.

Each unit can be reserved for the exclusive use of a host or group of hosts. See “[Restricting Host Access in Multiple-Bus Failover Mode](#)” on page 52.

### Assigning a Unit Number to a StorageSet

To assign a unit number to a storageset, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESET-NAME
```

For example:

To assign unit D102 to RAIDset R1, use the following command:

```
ADD UNIT D102 R1
```

### Assigning a Unit Number to a Single (JBOD) Disk

To assign a unit number to a single (JBOD) disk, use the following syntax:

```
ADD UNIT UNIT-NUMBER DISK-NAME
```

For example:

To assign unit D4 to DISK20300, use the following command:

```
ADD UNIT D4 DISK20300
```

### Assigning a Unit Number to a Partition

To assign a unit number to a partition, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESET-NAME PARTITION=PARTITION-NUMBER
```

For example:

To assign unit D100 to partition 3 of mirrorset mirr1, use the following command:

```
ADD UNIT D100 MIRR1 PARTITION=3
```

## Preferring Units

In multiple-bus failover mode, individual units can be preferred to a specific controller. For example, to prefer unit D102 to “this controller,” use the following command:

```
SET D102 PREFERRED_PATH=THIS
```

RESTART commands must be issued to both controllers for this command to take effect:

```
RESTART OTHER_CONTROLLER  
RESTART THIS_CONTROLLER
```

---

**Note:** The controllers need to restart together for the preferred settings to take effect. The `RESTART this_controller` command must be entered immediately after the `RESTART other_controller` command.

---

## Configuration Options

### Changing the CLI Prompt

To change the CLI prompt, enter a 1- to 16-character string as the new prompt, according to the following syntax:

```
SET THIS_CONTROLLER PROMPT = "NEW PROMPT"
```

If you are configuring dual-redundant controllers, also change the CLI prompt on the “other controller.” Use the following syntax:

```
SET OTHER_CONTROLLER PROMPT = "NEW PROMPT"
```

---

**Note:** It is suggested that the prompt name reflect some information about the controllers. For example, if the subsystem is the third one in a lab, name the top controller prompt, LAB3A and the bottom controller, LAB3B.

---

### Mirroring cache

To specify mirrored cache, use the following syntax:

```
SET THIS MIRRORED_CACHE
```

## Adding Disk Drives

If you add new disk drives to the subsystem, the disk drives must be added to the controllers' list of known devices:

- To add one new disk drive to the list of known devices, use the following syntax:

```
ADD DISK DISKNNNNN P T L
```

- To add several new disk drives to the list of known devices, enter the following command:

```
RUN CONFIG
```

## Adding a Disk Drive to the Spareset

The spareset is a collection of spare disk drives that are available to the controller should it need to replace a failed member of a RAIDset or mirrorset.

---

**Note:** This procedure assumes that the disks that you are adding to the spareset have already been added to the controller's list of known devices.

---

To add the disk drive to the controller's spareset list, use the following syntax:

```
ADD SPARESET DISKNNNNN
```

Repeat this step for each disk drive you want to add to the spareset:

For example:

The following example shows the syntax for adding DISK11300 and DISK21300 to the spareset.

```
ADD SPARESET DISK11300
```

```
ADD SPARESET DISK21300
```

## Removing a Disk Drive from the Spareset

You can delete disks in the spareset if you need to use them elsewhere in your subsystem.

1. Show the contents of the spareset entering the following command:

```
SHOW SPARESET
```

2. Delete the desired disk drive entering the following command:

```
DELETE SPARESET DISKNNNNNN
```

The RUN CONFIG command does not delete disks from the controllers' device table if a disk has been physically removed or replaced. In this case, you must use the command: `DELETE DISKNNNNNN`.

3. Verify the contents of the spareset by entering the following command:

```
SHOW SPARESET
```

## Enabling Autospare

With AUTOSPARE enabled on the failedset, any new disk drive that is inserted into the PTL location of a failed disk drive is automatically initialized and placed into the spareset. If initialization fails, the disk drive remains in the failedset until you manually delete it from the failedset.

To enable autospare, use the following command:

```
SET FAILEDSET AUTOSPARE
```

To disable autospare, use the following command:

```
SET FAILEDSET NOAUTOSPARE
```

During initialization, AUTOSPARE checks to see if the new disk drive contains metadata. Metadata is information the controller writes on the disk drive when the disk drive is configured into a storage set. Therefore, the presence of metadata indicates that the disk drive belongs to, or has been used by, a storage set. If the disk drive contains metadata, initialization stops.

---

**Note:** A new disk drive will not contain metadata but a repaired or reused disk drive might. To erase metadata from a disk drive, add it to the controller's list of devices, then set it to be nontransportable and initialize it.

---

## Deleting a Storageset

---

**Note:** If the storageset you are deleting is partitioned, you must delete each partitioned unit before you can delete the storageset.

---

1. Show the storageset's configuration:

```
SHOW STORAGESET-NAME
```

2. Delete the unit number that uses the storageset. Use the following command:

```
DELETE UNIT-NUMBER
```

3. Delete the storageset. Use the following command:

```
DELETE STORAGESET-NAME
```

4. Verify the configuration:

```
SHOW STORAGESET-NAME
```

## Changing Switches for a Storageset or Device

You can optimize a storageset or device at any time by changing the switches that are associated with it. Remember to update the storageset profile when changing its switches.

### Displaying the Current Switches

To display the current switches for a storageset or single-disk unit, enter a SHOW command, specifying the FULL switch:

```
SHOW STORAGESET-NAME
```

or

```
SHOW DEVICE-NAME
```

---

**Note:** FULL is not required when showing a particular device. It is used when showing all devices, for example, SHOW DEVICES FULL.

---

### Changing RAIDset and Mirrorset Switches

Use the SET storageset-name command to change the RAIDset and Mirrorset switches associated with an existing storageset.

For example, the following command changes the replacement policy for RAIDset RAID1 to BEST\_FIT:

```
SET RAID1 POLICY=BEST_FIT
```

## Changing Device Switches

Use the `SET device-name` command to change the device switches.

For example, to request a data transfer rate of 20 MHz for DISK10000:

```
SET DISK10000 TRANSFER_RATE_REQUESTED=20MHZ
```

## Changing Initialize Switches

The initialization switches cannot be changed without destroying the data on the storageset or device. These switches are integral to the formatting and can only be changed by reinitializing the storageset. Initializing a storageset is similar to formatting a disk drive; all data is destroyed during this procedure.

## Changing Unit Switches

Use the `SET unit-name` command to change the characteristics of a unit.

For example, the following command enables write protection for unit D100:

```
SET D100 WRITE_PROTECT
```

## Verifying Storage Configuration from Host

This section briefly describes how to verify that multiple paths exist to virtual disk units under Tru64 UNIX V5.1x.

After configuring units (virtual disks) through either the CLI or SWCC, access the new storage by using one of the following methods:

Issuing the following command to rescan the bus:

```
# hwmgr -scan scsi
# dsfmgr -k
```

or

Restart the host

After the host restarts, verify that the disk is correctly presented to the host. The command to use consists of the following syntax:

```
# hwmgr -view dev
```

The disk information returned is shown in the following table:

HWDI: Device Name	Mfg	Model	Location
123: /dev/disk/dsk56c	DEC	HSG80C CL	bus-5-targ-0-lun-0
124: /dev/disk/dsk57c	DEC	HSG80	bus-6-targ-0-lun-1
125: /dev/disk/dsk58c	DEC	HSG80	bus-5-targ-0-lun-2
126: /dev/disk/dsk59c	DEC	HSG80	bus-5-targ-0-lun-60
127: /dev/disk/dsk60c	DEC	HSGCCL	bus-5-targ-10-lun-0
128: /dev/disk/dsk61c	DEC	HSG80	bus-5-targ-10-lun-60
129: /dev/disk/dsk62c	DEC	HSG80	bus-5-targ-10-lun-61
130: /dev/disk/dsk63c	DEC	HSG80	bus-5-targ-10-lun-62



# Using CLI for Configuration

6

This chapter presents an example of how to configure a storage subsystem using the Command Line Interpreter (CLI). The CLI configuration example shown assumes:

- A normal, new controller pair, which includes:
  - NODE ID set
  - No previous failover mode
  - No previous topology set
- Full array with no expansion cabinet
- PCMCIA cards installed in both controllers

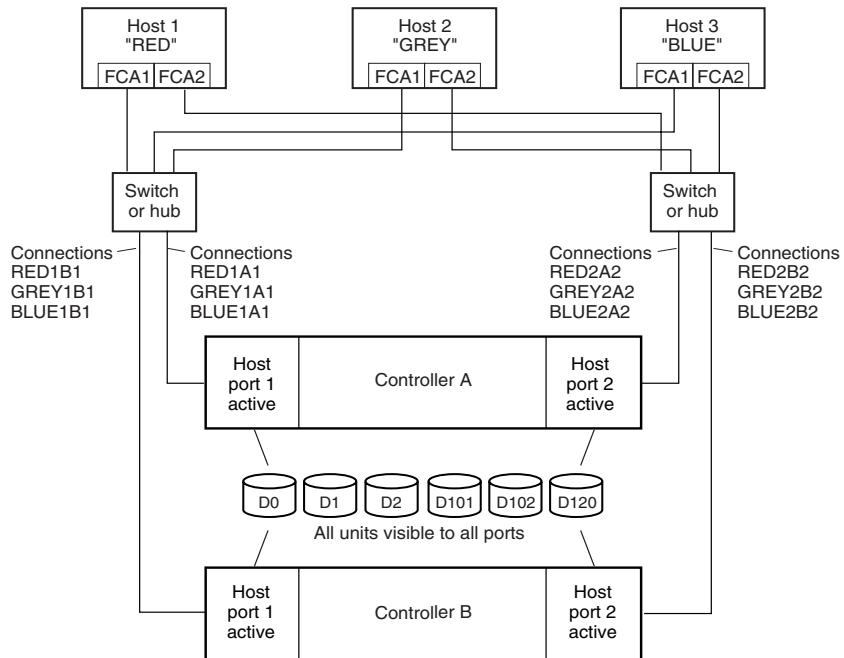
A storage subsystem example is shown in [Figure 40](#). The example system contains three non-clustered TRU64\_UNIX hosts, as shown in [Figure 43](#). From the hosts' point of view, each host will have four paths to its own virtual disks. The resulting virtual system, from the host's point of view, is shown in [Figure 43](#).

Figure 40 shows an example storage system map for the BA370 enclosure. Details on building your own map are described in Chapter 2. Templates to help you build your storage map are supplied in Appendix A.

Port											Targets
	1	2	3	4	5	6					Targets
Power Supply	D2 S2 DISK10300	D2 S2 DISK20300	D2 S2 DISK30300	D2 S2 DISK40300	D2 D101 DISK50300	spareset member DISK60300	Power Supply				3
Power Supply	D0 S1 MI DISK10200	D0 S1 M1 DISK20200	D0 S1 M2 DISK30200	D0 S1 M2 DISK40200	D1 M3 DISK50200	D1 M3 DISK60200	Power Supply				2
Power Supply	D120 R2 DISK10100	D120 R2 DISK20100	D120 R2 DISK30100	D120 R2 DISK40100	D120 R2 DISK50100	D120 R2 DISK60100	Power Supply				1
Power Supply	D102 R1 DISK10000	D102 R1 DISK20000	D102 R1 DISK30000	D102 R1 DISK40000	D102 R1 DISK50000	D102 R1 DISK60000	Power Supply				0

**Figure 40: Example storage map for the BA370 Enclosure**

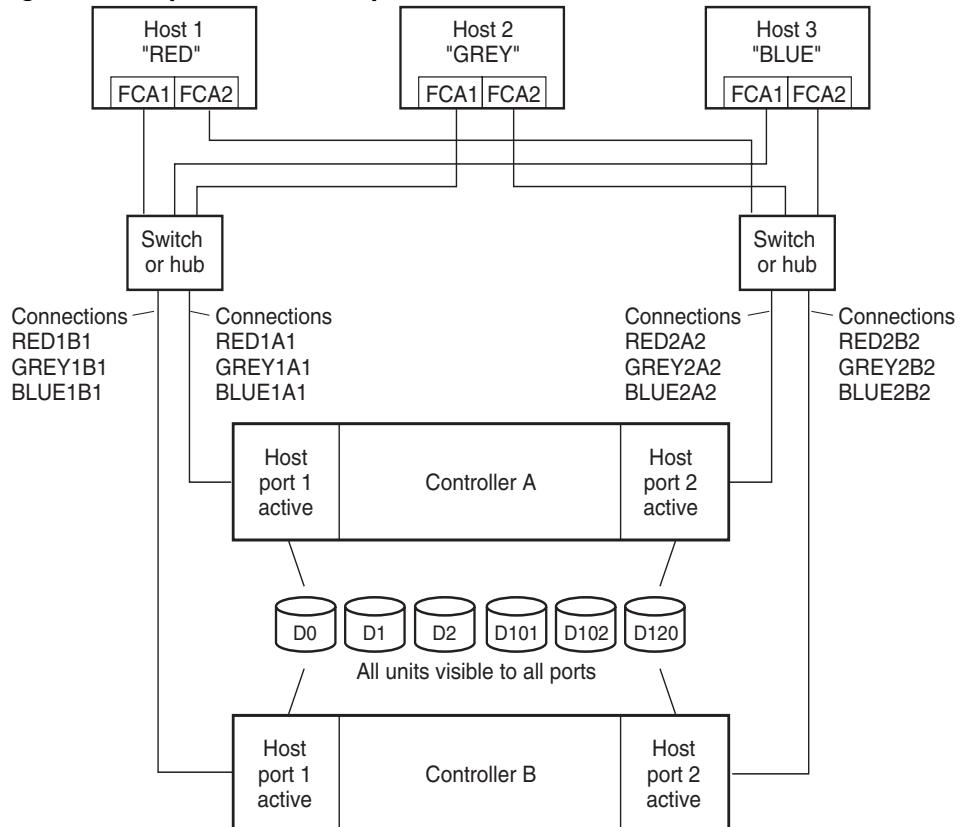
Figure 41 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host BLUE can be done by enabling only the connections to host BLUE. At least two connections must be enabled for multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. The example system, shown in Figure 43, contains three non-clustered TRU64\_UNIX hosts. Port 1 link is separate from port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another). Therefore, each adapter has two connections.



NOTE: FCA = Fibre Channel Adapter

CXO7109B

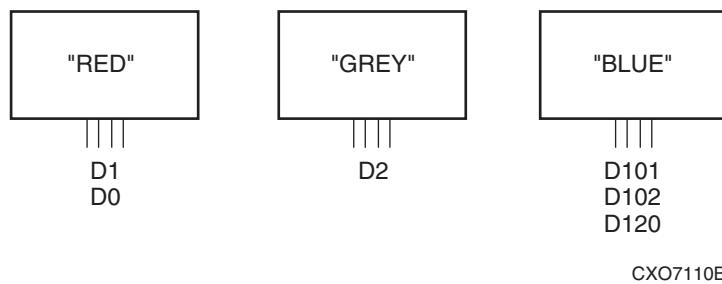
**Figure 41: Example, three non-clustered host systems**

**Figure 42: Representative multiple-bus failover**

NOTE: FCA = Fibre Channel Adapter

CX07547B

**Figure 43** represents units that are logical or virtual disks comprised of storage sets configured from physical disks.



**Figure 43: Example, logical or virtual disks comprised of storage sets**

## CLI Configuration Example

Text conventions used in this example are listed below:

- Text in *italics* indicates an action you take.
- Text in **THIS FORMAT**, indicates a command you type. Be certain to press **Enter** after each command.
- Text enclosed within a box, indicates information that is displayed by the CLI interpreter.

---

**Note:** "This" controller is top controller (A).

---

*Plug serial cable from maintenance terminal into top controller.*

```
CLEAR CLI  
SET MULTIBUS COPY=THIS  
CLEAR CLI
```

```
SET THIS SCSI_VERSION=SCSI-3
SET THIS PORT_1_TOPOLOGY=FABRIC
SET THIS PORT_2_TOPOLOGY=FABRIC
SET OTHER PORT_1_TOPOLOGY=FABRIC
SET OTHER PORT_2_TOPOLOGY=FABRIC
SET THIS ALLOCATION_CLASS=0
RESTART OTHER
RESTART THIS
SET THIS TIME=10-Mar-2001:12:30:34
RUN FRUTIL
```

Do you intend to replace this controller's cache  
battery? Y/N [Y]

Y

Plug serial cable from maintenance terminal into bottom controller.

---

**Note:** Bottom controller (B) becomes “this” controller.

---

```
RUN FRUTIL
```

Do you intend to replace this controller's cache  
battery? Y/N [Y]

Y

```
SET THIS MIRRORED_CACHE
```

---

**Note:** This command causes the controllers to restart.

---

```
SET THIS PROMPT="BTVS BOTTOM"
SET OTHER PROMPT="BTVS TOP"
SHOW THIS
SHOW OTHER
```

Plug in the Fibre Channel cable from the first adapter in host “RED.”

```
SHOW CONNECTIONS
```

```

RENAME !NEWCON00 RED1B1
SET RED1B1 OPERATING_SYSTEM=TRU64_UNIX
RENAME !NEWCON01 RED1A1
SET RED1A1 OPERATING_SYSTEM=TRU64_UNIX
SHOW CONNECTIONS

```

---

**Note:** Connection table sorts alphabetically.

---

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
RED1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
RED1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		

Mark or tag both ends of Fibre Channel cables.

Plug in the Fibre Channel cable from the second adapter in host “RED.”

SHOW CONNECTIONS

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON02	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
!NEWCON03	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX				ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
RED1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
...						

```

RENAME !NEWCON02 RED2B2
SET RED2B2 OPERATING_SYSTEM=TRU64_UNIX
RENAME !NEWCON02 RED2A2
SET RED2A2 OPERATING_SYSTEM=TRU64_UNIX
SHOW CONNECTIONS

```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
RED1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

Mark or tag both ends of Fibre Channel cables

Repeat this process to add connections from the other two hosts. The resulting connection table should appear similar to the following:

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
GREY1A1	TRU64_UNI	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
GREY1B1	TRU64_UNI	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
GREY2A2	TRU64_UNI	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
GREY2B2	TRU64_UNI	THIS	2	XXXXXX	OL this	0

(Continued)

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BLUE1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BLUE1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BLUE2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BLUE2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
RED2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

```
RUN CONFIG
ADD RAIDSET R1 DISK10000 DISK20000 DISK30000 DISK40000 DISK50000
DISK60000
INITIALIZE R1
ADD UNIT D102 R1 DISABLE_ACCESS_PATH=ALL
SET D102 ENABLE_ACCESS_PATH=(RED1A1, RED1B1, RED2A2, RED2B2)
ADD RAIDSET R2 DISK10100 DISK20100 DISK30100 DISK40100 DISK50100
DISK60100
INITIALIZE R2
ADD UNIT D120 R2 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE1B1, BLUE2A2,
BLUE2B2)
ADD MIRRORSET MI DISK10200 DISK20200
ADD MIRRORSET M2 DISK30200 DISK40200
ADD STRIPESET S1 M1 M2
INITIALIZE S1
ADD UNIT D0 S1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=(GREY1A1, GREY1B1, GREY2A2, GREY2B2)
ADD MIRRORSET M3 DISK50200 DISK60200
INITIALIZE M3
ADD UNIT D1 M3 DISABLE_ACCESS_PATH=ALL
SET D1 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE1B1, BLUE2A2, BLUE2B2)
ADD STRIPESET S2 DISK10300 DISK20300 DISK30300 DISK40300
INITIALIZE S2
ADD UNIT D2 S2 DISABLE_ACCESS_PATH=ALL
SET D2 ENABLE_ACCESS_PATH=(GREY1A1, GREY1B1, GREY2A2, GREY2B2)
INITIALIZE DISK50300
ADD UNIT D101 DISK50300 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE1B1, BLUE2A2,
BLUE2B2)
ADD SPARESET DISK60300
SHOW UNITS FULL
```

# Backing Up, Cloning, and Moving Data

This chapter includes the following topics:

- [Backing Up Subsystem Configurations](#), page 188
- [Creating Clones for Backup](#), page 189
- [Moving Storagesets](#), page 193

## Backing Up Subsystem Configurations

The controller stores information about the subsystem configuration in its nonvolatile memory. This information could be lost if the controller fails or when you replace a module in the subsystem.

Use the following command to produce a display that shows if the save configuration feature is active and which devices are being used to store the configuration.

```
SHOW THIS_CONTROLLER FULL
```

The resulting display includes a line that indicates status and how many devices have copies of the configuration. The last line shows on how many devices the configuration is backed up.

---

**Note:** Do not use `SAVE_CONFIGURATION` in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

---

The `SHOW DEVICES FULL` command shows which disk drives are set up to back up the configuration. The syntax for this command is shown below:

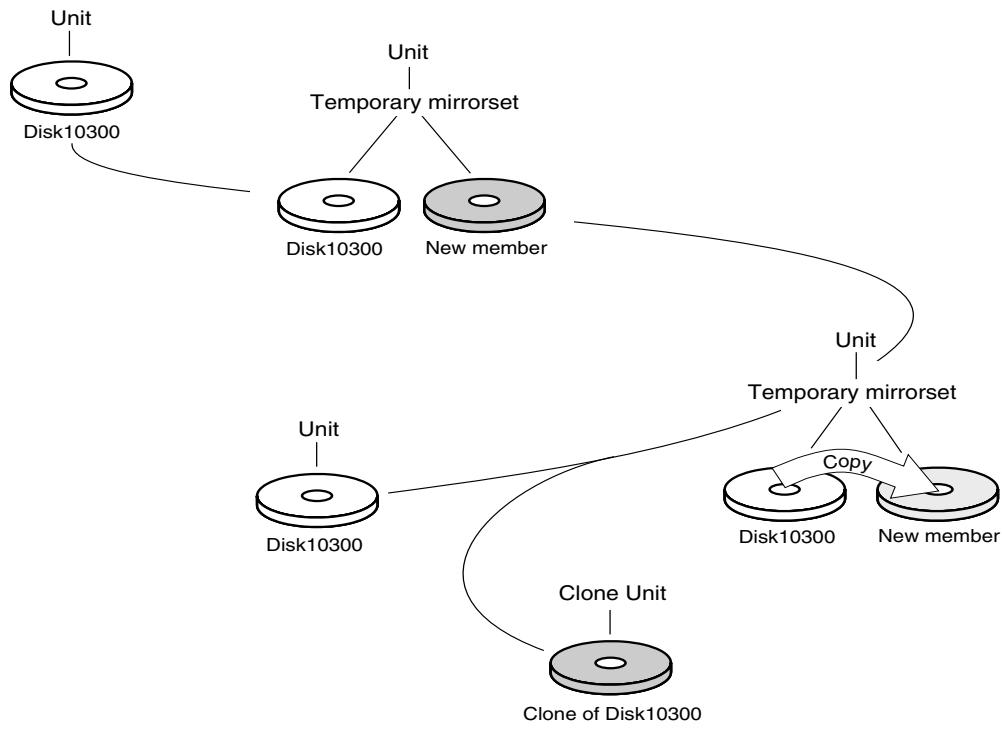
```
SHOW DEVICES FULL
```

## Creating Clones for Backup

Use the CLONE utility to duplicate the data on any unpartitioned single-disk unit, stripeset, mirrorset, or striped mirrorset in preparation for backup. When the cloning operation is complete, you can back up the clones rather than the storage set or single-disk unit, which can continue to service its I/O load. When you are cloning a mirrorset, CLONE does not need to create a temporary mirrorset. Instead, it adds a temporary member to the mirrorset and copies the data onto this new member.

The CLONE utility creates a temporary, two-member mirrorset for each member in a single-disk unit or stripeset. Each temporary mirrorset contains one disk drive from the unit you are cloning and one disk drive onto which CLONE copies the data. During the copy operation, the unit remains online and active so that the clones contain the most up-to-date data.

After the CLONE utility copies the data from the members to the clones, it restores the unit to its original configuration and creates a clone unit you can back up. The CLONE utility uses the steps shown in [Figure 44](#) to duplicate each member of a unit.



**Figure 44: CLONE utility steps for duplicating unit members**

Use the following steps to clone a single-disk unit, stripeset, or mirrorset:

1. Establish a connection to the controller that accesses the unit you want to clone.
2. Start CLONE using the following command:  
`RUN CLONE`
3. When prompted, enter the unit number of the unit you want to clone.
4. When prompted, enter a unit number for the clone unit that CLONE will create.
5. When prompted, indicate how you would like the clone unit to be brought online: either automatically or only after your approval.
6. When prompted, enter the disk drives you want to use for the clone units.
7. Back up the clone unit.

The following example shows the commands you would use to clone storage unit D6. The clone command terminates after it creates storage unit D33, a clone or copy of D6.

```
RUN CLONE
CLONE LOCAL PROGRAM INVOKED
UNITS AVAILABLE FOR CLONING:
98
ENTER UNIT TO CLONE? 98
CLONE WILL CREATE A NEW UNIT WHICH IS A COPY OF UNIT 98.
ENTER THE UNIT NUMBER WHICH YOU WANT ASSIGNED TO THE NEW UNIT?
99
THE NEW UNIT MAY BE ADDED USING ONE OF THE FOLLOWING METHODS:
1. CLONE WILL PAUSE AFTER ALL MEMBERS HAVE BEEN COPIED. THE USER
MUST THEN PRESS RETURN TO CAUSE THE NEW UNIT TO BE ADDED.
2. AFTER ALL MEMBERS HAVE BEEN COPIED, THE UNIT WILL BE ADDED
AUTOMATICALLY.

UNDER WHICH ABOVE METHOD SHOULD THE NEW UNIT BE ADDED[ ]?1
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20200 (SIZE=832317)
DISK20300 (SIZE=832317)
USE AVAILABLE DEVICE DISK20200(SIZE=832317) FOR MEMBER
DISK10300(SIZE=832317) (Y,N) [Y]? Y
MIRROR DISK10300 C_MA
SET C_MA NOPOLICY
SET C_MA MEMBERS=2
SET C_MA REPLACE=DISK20200
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20300 (SIZE=832317)
```

```
USE AVAILABLE DEVICE DISK20300(SIZE=832317) FOR MEMBER
DISK10000(SIZE=832317) (Y,N) [Y]? Y
MIRROR DISK10000 C_MB
SET C_MB NOPOLICY
SET C_MB MEMBERS=2
SET C_MB REPLACE=DISK20300
COPY IN PROGRESS FOR EACH NEW MEMBER. PLEASE BE PATIENT...
.
.
COPY FROM DISK10300 TO DISK20200 IS 100% COMPLETE
COPY FROM DISK10000 TO DISK20300 IS 100% COMPLETE
PRESS RETURN WHEN YOU WANT THE NEW UNIT TO BE CREATED
REDUCE DISK20200 DISK20300
UNMIRROR DISK10300
UNMIRROR DISK10000
ADD MIRRORSET C_MA          DISK20200
ADD MIRRORSET C_MB          DISK20300
ADD STRIPESET C_ST1 C_MA C_MB
INIT C_ST1      NODESTROY
ADD UNIT D99 C_ST1
D99 HAS BEEN CREATED. IT IS A CLONE OF D98.
CLONE - NORMAL TERMINATION
```

## Moving StorageSets

You can move a storageSet from one subsystem to another without destroying its data. You also can follow the steps in this section to move a storageSet to a new location within the same subsystem.



**Caution:** Move only normal storageSets. Do not move storageSets that are reconstructing or reduced, or data corruption will result.

See the release notes for the version of your controller software for information on which drives can be supported.



**Caution:** Never initialize any container or this procedure will not protect data in the storageSet.

Use the following procedure to move a storageSet, while maintaining the data the storageSet contains:

1. Show the details for the storageSet you want to move. Use the following command:

```
SHOW STORAGESET-NAME
```

2. Label each member with its name and PTL location.

If you do not have a storageSet map for your subsystem, you can enter the LOCATE command for each member to find its PTL location. Use the following command:

```
LOCATE DISK-NAME
```

To cancel the locate command, enter the following:

```
LOCATE CANCEL
```

3. Delete the unit number shown in the “Used by” column of the SHOW storageSet-name command. Use the following syntax:

```
DELETE UNIT-NUMBER
```

4. Delete the storageSet shown in the “Name” column of the SHOW storageSet-name command. Use the following syntax:

```
DELETE STORAGESET-NAME
```

5. Delete each disk drive, one at a time, that the storageset contained. Use the following syntax:

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

6. Remove the disk drives and move them to their new PTL locations.
7. Again add each disk drive to the controller's list of valid devices. Use the following syntax:

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

8. Recreate the storageset by adding its name to the controller's list of valid storagesets and by specifying the disk drives it contains. (Although you have to recreate the storageset from its original disks, you do not have to add the storagesets in their original order.) Use the following syntax to recreate the storageset:

```
ADD STORAGESET-NAME DISK-NAME DISK-NAME
```

9. Represent the storageset to the host by giving it a unit number the host can recognize. You can use the original unit number or create a new one. Use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESET-NAME
```

The following example moves unit D100 to another cabinet. D100 is the RAIDset RAID99 that consists of members DISK10000, DISK20000, and DISK10100.

Old cabinet

```
DELETE D100
```

```
DELETE RAID99
```

```
DELETE DISK10000
```

```
DELETE DISK10100
```

```
DELETE DISK20000
```

```
DELETE DISK20100
```

New cabinet

```
ADD DISK DISK10000
ADD DISK DISK10100
ADD DISK DISK20000
ADD DISK DISK20100
ADD RAIDSET RAID99 DISK10000 DISK10100 DISK20000 DISK20100
ADD UNIT D100 RAID99
```



# Subsystem Profile Templates

A

This appendix contains storageset profiles to copy and use to create your profiles. It also contains an enclosure template to use to help keep track of the location of devices and storagesets in your shelves. Four (4) templates will be needed for the subsystem.

---

**Note:** The storage map templates for the Model 4310R and Model 4214R or 4314R reflect the physical location of the disk enclosures in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

---

- [Storageset Profile](#), page 198
- [Storage Map Template 1 for the BA370 Enclosure](#), page 200
- [Storage Map Template 2 for the Second BA370 Enclosure](#), page 201
- [Storage Map Template 3 for the TSsingle-Bushird BA370 Enclosure](#), page 202
- [Storage Map Template 4 for the Model 4214R Disk Enclosure](#), page 203
- [Storage Map Template 5 for the Model 4254 Disk Enclosure](#), page 205
- [Storage Map Template 6 for the Model 4310R Disk Enclosure](#), page 207
- [Storage Map Template 7 for the Model 4350R Disk Enclosure](#), page 209
- [Storage Map Template 8 for the Model 4314R Disk Enclosure](#), page 210
- [Storage Map Template 9 for the Model 4354R Disk Enclosure](#), page 212

## Storageset Profile

Type of Storageset:

<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD
Mirrorset				

**Storageset Name****Disk Drives****Unit Number**

Partitions:


RAIDset Switches:

<input type="checkbox"/> Normal (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Best performance (default)
<input type="checkbox"/> Fast	<input type="checkbox"/> Yes, missing:	<input type="checkbox"/> Best fit
<input type="checkbox"/> None		

Mirrorset Switches:

<input type="checkbox"/> Best performance (default)	<input type="checkbox"/> Normal (default)	<input type="checkbox"/> Least busy (default)
<input type="checkbox"/> Best fit	<input type="checkbox"/> Fast	<input type="checkbox"/> Round robin
<input type="checkbox"/> None	<input type="checkbox"/> Disk drive:	

Initialize Switches:

<input type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Destroy (default)
<input type="checkbox"/> 64 blocks	<input type="checkbox"/> Yes	<input type="checkbox"/> Retain
<input type="checkbox"/> 128 blocks		
<input type="checkbox"/> 256 blocks		
<input type="checkbox"/> Other:		

### Unit Switches:

Read caching \_\_\_\_\_

Read-ahead caching \_\_\_\_\_

Write-back caching \_\_\_\_\_

Write-through caching \_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

\_\_\_\_\_

## Storage Map Template 1 for the BA370 Enclosure

Use this template for:

- BA370 single-enclosure subsystems
- first enclosure of multiple BA370 enclosure subsystems

	Port						
	1	2	3	4	5	6	
Power Supply							Power Supply
	D10300	D20300	D30300	D40300	D50300	D60300	
Power Supply							Power Supply
	D10200	D20200	D30200	D40200	D50200	D60200	
Power Supply							Power Supply
	D10100	D20100	D30100	D40100	D50100	D60100	
Power Supply							Power Supply
	D10000	D20000	D30000	D40000	D50000	D60000	

## Storage Map Template 2 for the Second BA370 Enclosure

Use this template for the second enclosure of multiple BA370 enclosure subsystems.

	Port						
	1	2	3	4	5	6	
Power Supply							Power Supply
	D11100	D21100	D31100	D41100	D51100	D61100	
Power Supply							Power Supply
	D11000	D21000	D31000	D41000	D51000	D61000	
Power Supply							Power Supply
	D10900	D20900	D30900	D40900	D50900	D60900	
Power Supply							Power Supply
	D10800	D20800	D30800	D40800	D50800	D60800	

## Storage Map Template 3 for the TSsingle-Bushird BA370 Enclosure

Use this template for the third enclosure of multiple BA370 enclosure subsystems.

	Port						
	1	2	3	4	5	6	
Power Supply							Power Supply
	D11500	D21500	D31500	D41500	D51500	D61500	
Power Supply							Power Supply
	D11400	D21400	D31400	D41400	D51400	D61400	
Power Supply							Power Supply
	D11300	D21300	D31300	D41300	D51300	D61300	
Power Supply							Power Supply
	D11200	D21200	D31200	D41200	D51200	D61200	

## Storage Map Template 4 for the Model 4214R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4214R disk enclosure (single-bus). You can have up to six Model 4214R disk enclosures per controller shelf.

Model 4214R Disk Enclosure Shelf 1 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	1	1	1	1	1
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500
Model 4214R Disk Enclosure Shelf 2 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	1	1	1	1	1
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

(Continued)

Model 4214R Disk Enclosure Shelf 3 (Single-Bus)																
Bay	1	2	3	4	5	6	7	8	9	1	0	1	1	2	3	4
SCSI ID	0 0	0 1	0 2	0 3	0 4	0 5	0 8	0 9	1 0	1 1	1 2	1 3	1 4	1 5		
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500		

## Storage Map Template 5 for the Model 4254 Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4254 disk enclosure (dual-bus). You can have up to three Model 4254 disk enclosures per controller shelf.

Model 4254 Disk Enclosure Shelf 1 (Dual-Bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	1	1	1	1	1
SCSI ID	0 0	0 1	0 2	0 3	0 4	0 5	0 8	0 0	0 1	0 2	0 3	0 4	0 5	0 8
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800

Model 4254 Disk Enclosure Shelf 2 (Dual-Bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	1	1	1	1	1
SCSI ID	0 0	0 1	0 2	0 3	0 4	0 5	0 8	0 0	0 1	0 2	0 3	0 4	0 5	0 8
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

(Continued)

Model 4254 Disk Enclosure Shelf 3 (Dual-Bus)														
Bay	Bus A							Bus B						
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

## Storage Map Template 6 for the Model 4310R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4310R disk enclosure (single-bus). You can have up to six Model 4310R disk enclosures per controller shelf.

Model 4310R Disk Enclosure Shelf 6 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200
Model 4310R Disk Enclosure Shelf 5 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200
Model 4310R Disk Enclosure Shelf 4 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Model 4310R Disk Enclosure Shelf 1 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk11000	Disk11100	Disk11200
Model 4310R Disk Enclosure Shelf 2 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk21000	Disk21100	Disk21200
Model 4310R Disk Enclosure Shelf 3 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

## Storage Map Template 7 for the Model 4350R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4350R disk enclosure (single-bus). You can have up to three Model 4350R disk enclosures per controller shelf.

Model 4350R Disk Enclosure Shelf 6 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200
Model 4350R Disk Enclosure Shelf 5 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200
Model 4350R Disk Enclosure Shelf 4 (Single-Bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

## Storage Map Template 8 for the Model 4314R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4314R disk enclosure. You can have a maximum of six Model 4314R disk enclosures with each Model 2200 controller enclosure.

Model 4314R Disk Enclosure Shelf 6 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500
Model 4314R Disk Enclosure Shelf 5 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51500	Disk51500
Model 4314R Disk Enclosure Shelf 4 (Single-Bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	Disk41400	Disk41500

continued from previous page

**Model 4314R Disk Enclosure Shelf 1 (Single-Bus)**

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

**Model 4314R Disk Enclosure Shelf 2 (Single-Bus)**

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

**Model 4314R Disk Enclosure Shelf 3 (Single-Bus)**

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

## Storage Map Template 9 for the Model 4354R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4354R disk enclosure (dual-bus). You can have up to three Model 4354R disk enclosures per controller shelf.

Model 4354R Disk Enclosure Shelf 1 (Dual-Bus)														
	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk21000	Disk22000	Disk23000	Disk24000	Disk25000	Disk20800
Model 4354R Disk Enclosure Shelf 2 (Dual-Bus)														
	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800
Model 4354R Disk Enclosure Shelf 3 (Dual-Bus)														
	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

# Installing, Configuring, and Removing the Client

B

The following information is included in this appendix:

- [Why Install the Client?](#), page 214
- [Before You Install the Client](#), page 215
- [Installing the Client](#), page 216
- [Installing the Integration Patch](#), page 217
- [Troubleshooting Client Installation](#), page 219
- [Adding Storage Subsystem and its Host to Navigation Tree](#), page 221
- [Removing Command Console Client](#), page 223
- [Where to Find Additional Information](#), page 224

## Why Install the Client?

The Client monitors and manages a storage subsystem by performing the following tasks:

- Create mirrored device group (RAID 1)
- Create striped device group (RAID 0)
- Create striped mirrored device group (RAID 0+1)
- Create striped parity device group (3/5)
- Create an individual device (JBOD)
- Monitor many subsystems at once
- Set up pager notification

## Before You Install the Client

1. Verify you are logged into an account that is a member of the administrator group.
2. Check the software product description that came with the software for a list of supported hardware.
3. Verify that you have the SNMP service installed on the computer. SNMP must be installed on the computer for this software to work properly. The Client software uses SNMP to receive traps from the Agent. The SNMP service is available on the Windows NT/Windows 2000/Windows Server 2003 (32-bit) installation CD-ROM. To verify that you have the SNMP service:
  - For Windows NT, double-click Services in *Start > Settings > Control Panel*. The entry for SNMP is shown in this window. If you install the SNMP service and you already have Windows NT Service Pack 6A on the computer, reinstall the service pack after installing the SNMP service.
  - For Windows 2000, click *Start > Settings > Control Panel > Administrative Tools > Component Services*. The entry for SNMP is shown in the Component Services window.
  - For Windows 2003 (32-bit), click *Start > Settings > Control Panel > Administrative Tools > Component Services*. The entry for SNMP is shown in the Component Services window.
4. Read the release notes.
5. Read “[Installing the Integration Patch](#)” on page 217 in this appendix.
6. If you have the Command Console Client open, exit the Command Console Client.
7. If you have Command Console Client version 1.1b or earlier, remove the program with the Windows Add/Remove Programs utility.
8. If you have a previous version of Command Console, you can save the Navigation Tree configuration by copying the *SWCC2.MDB* file to another directory. After you have installed the product, move *SWCC2.MDB* to the directory to which you installed SWCC.
9. Install the HS-Series Agent. For more information, see Chapter 4.

## Installing the Client

The following restriction should be observed when installing SWCC on Windows NT 4.0 Workstations.

If you select all of the applets during installation, the installation will fail on the HSG60 applet and again on one of the HSG80 applets. The workaround is to install all of the applets you want except for the HSG60 applet and the HSG80 ACS 8.5 applet. You can then return to the setup program and install the one that you need.

1. In a SAN environment where you would need both HSG60 and HSG80 subsystems, StorageWorks recommends you install both, but one at a time. This problem is not seen under Windows NT 4.0 Server.
2. Insert the CD-ROM into a computer running Windows 2000 with Service Pack 3 or 4, or Windows NT 4.0 (Intel) with Service Pack 6.0A, or Windows Server 2003 (32-bit).
3. A dialog box should automatically appear.
4. One of the items in the dialog box should say “SWCC Client Software” and has a button that says “INSTALL” next to it. Click on the button to start the SWCC client installation procedure.
5. Select *HSG80 Controller for ACS87 or newer* menu option to properly install SWCC client, and click **Next**.

If this method does not work, go to the \windows directory on the CD-ROM and run the *setup.exe* program.

---

**Note:** If the computer does not find a previous installation, it will install the SWCC Navigation Window and the CLI Window.

---

6. Follow the instructions on the screen. After you install the software, the Asynchronous Event Service (AES) starts. AES is a service that runs in the background. It collects and passes traps from the subsystems to the Navigation Tree and to individual pagers (for example, to show that a disk has failed). AES needs to be running for the client system to receive updates.

---

**Note:** For more information on AES, see *HP StorageWorks Command Console Version 2.5, User Guide*.

---

## Installing the Integration Patch

The integration patch determines which version of firmware the controller is using and launches the appropriate StorageWorks Command Console (SWCC) Storage Window within Insight Manager (CIM) version 4.23.

### Should I Install the Integration Patch?

Install this patch if your HSG80 controller uses ACS 8.7 or later. This patch enables you to use the controller's SWCC Storage Window within CIM to monitor and manage the controller.

### Performing the Integration Patch Installation

Perform the following steps to install the integration patch.

1. Verify that you have installed the HSG80 Storage Window for ACS 8.7 or later in the Add/Remove Programs in the Windows Control Panel. The HSG80 Storage Window for ACS 8.7 or later is needed to display the correct Storage Window for your version of the firmware.
2. Verify that you have installed HSG80 Storage Window version 2.1 in the Add/Remove Programs (StorageWorks HSG80 V2.1) in the Windows Control Panel. The HSG80 Storage Window version 2.1 is required to run the integration patch.
3. Verify that you have installed CIM version 4.23.
4. Install the integration patch from the Solution Software CD-ROM by double-clicking on setup.exe in the following directory:

\SWCC\Client\HSG80shim

The patch is installed in the same location as the original SWCC installation.



**Caution:** Do not remove the HSG80 Client from your computer. If you remove the HSG80 Client, you will no longer be able to access its Storage Window.

---



**Caution:** If you remove the integration patch, HSG80 Storage Window V2.1 will no longer work and you will need to reinstall HSG80 Storage Window V2.1. The integration patch uses some of the same files as the HSG80 Storage Window V2.1.

---

## Integrating Controller's SWCC Storage Window with CIM

You can open the controller's Storage Window from within the Windows-based CIM version 4.23 by doing the following:

1. Verify that you have installed the following by looking in Add/Remove Programs in Control Panel:
  - The HSG80 Storage Window for ACS 8.6 or later (Required to open the correct Storage Window for your firmware).
  - The HSG80 Storage Window version 2.1 (StorageWorks HSG80 V2.1) The CIM integration patch uses files in this program.CIM version 4.23.
  - CIM integration patch (HSG80 Insight Manager Shim).
2. Verify that you have installed the CIM Agent and the StorageWorks Command Console HS-Series Agent on the same computer.
3. Add the name of the client system that has CIM to the Agent's list of client system entries and choose SNMP as a notification scheme.
4. Open Insight Manager.
5. To open the Server window, click on the device you want to observe in the CIM Navigation window.
6. Click on the Mass Storage button in the Server window. The CIM Navigation Tree is displayed.
7. Click on the + symbol next to RAID Storage System. The Navigation Tree expands to display a listing called Storage System Information.
8. Double-click Storage System Information. You are given the status of the system.
9. Click Launch. The controller's Storage Window is displayed.

## Finding the Controller's Storage Window

If you installed Insight Manager before SWCC, Insight Manager will be unable to find the controller's Storage Window. To find the controller's Storage Window, perform the following procedure:

1. Double-click the Insight Agents icon (Start > Settings > Control Panel). A window appears showing you the active and inactive Agents under the Services tab.
2. Highlight the entry for Fibre Array Information and click Add. The Fibre Array Information entry is moved from Inactive Agents to Active Agents.

## Troubleshooting Client Installation

This section provides information on how to resolve some of the problems that may appear when installing the Client software:

- Invalid Network Port Assignments During Installation
- “There is no disk in the drive” Message

### Invalid Network Port Assignments During Installation

SWCC Clients and Agents communicate by using sockets. The SWCC installation attempts to add entries into each system list of services (services file or for UCX, the local services database). If the SWCC installation finds an entry in the local services file with the same name as the one it wants to add, it assumes the one in the file is correct.

The SWCC installation may display a message, stating that it cannot upgrade the services file. This happens if it finds an entry in the local services file with the same number as the one it wants to add, but with a different name. In that case, appropriate port numbers must be obtained for the network and added manually to the services file.

There are two default port numbers, one for Command Console (4998) and the other for the device-specific Agent and Client software, such as the Fibre Channel Interconnect Client and Agent (4989). There are two exceptions. The following software has two default port numbers:

- The KZPCC Agent and Client, (4991 and 4985)
- The RA200 Agent and Client, (4997 and 4995)

If the Network Information Services (NIS) are being used to provide named port lookup services, contact the network administrator to add the correct ports.

The following shows how the network port assignments appear in the services file:

spgui	4998/tc p	#Command Console
ccdevmgt	4993/tc p	#Device Management Client and Agent
kzpcccconnectport	4991/tc p	#KZPCC Client and Agent
kzpccdiscoveryport	4985/tc p	#KZPCC Client and Agent
ccfabric	4989/tc p	#Fibre Channel Interconnect Agent
spagent	4999/tc p	#HS-Series Client and Agent
spagent3	4994/tc p	#HSZ22 Client and Agent
ccagent	4997/tc p	#RA200 Client and Agent
spagent2	4995/tc p	#RA200 Client and Agent

## **"There is no disk in the drive" Message**

When you install the Command Console Client, the software checks the shortcuts on the desktop and in the Start menu. The installation will check the shortcuts of all users for that computer, even if they are not currently logged on. You may receive an error message if any of these shortcuts point to empty floppy drives, empty CD-ROM drives, or missing removable disks. Do one of the following:

- Ignore the error message by clicking **Ignore**.
- Replace the removable disks, and place a disk in the floppy drive and a CD-ROM in the CD-ROM drive. Then, click **Retry**.

## Adding Storage Subsystem and its Host to Navigation Tree

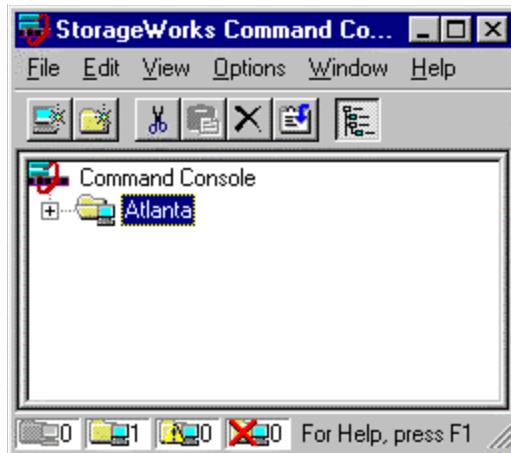
The Navigation Tree enables you to manage storage over the network by using the Storage Window. If you plan to use pager notification, you must add the storage subsystem to the Navigation Tree.

1. Verify that you have properly installed and configured the HS-Series Agent on the storage subsystem host.
2. Click *Start > Programs > Command Console > StorageWorks Command Console*. Client displays the Navigation Window. The Navigation Window lets you monitor and manage many storage subsystems over the network.



**Figure 45: Navigation Window**

3. Click *File > Add System*. The Add System window appears.
4. Type the host name or its TCP/IP address and click *Apply*.
5. Click *Close*.



**Figure 46: Navigation window showing storage host system "Atlanta"**

6. Click the plus sign to expand the host icon. When expanded, the Navigation Window displays an icon for the storage subsystem. To access the Storage Window for the subsystem, double-click the Storage Window icon.



**Figure 47: Navigation window showing expanded "Atlanta" host icon**

**Note:** You can create virtual disks by using the Storage Window. For more information on the Storage Window, refer to *HP StorageWorks Command Console Version 2.5, User Guide*.

---

## Removing Command Console Client

Before you remove the Command Console Client (CCL) from the computer, remove AES. This will prevent the system from reporting that a service failed to start every time the system is restarted. Steps 2 through 5 describe how to remove the CCL.

---

**Note:** When you remove the CCL, the *SWCC2.MDB* file is deleted. This file contains the Navigation Tree configuration. If you want to save this information, move the file to another directory.

---

1. Click *Start > Programs > Command Prompt* and change to the directory in which you installed the CCL.

2. Enter the following command:

```
C:\Program Files\Compaq\SWCC> AsyncEventService -remove
```

3. Do one of the following:

- On Windows NT 4.0, click *Start > Settings > Control Panel*, and double-click the Add/Remove Programs icon in the Control Panel. The Add/Remove Program Properties window appears.
- On Windows 2000, click *Start > Settings > Control Panel > Add/Remove Programs*. The Add/Remove Program window appears.

4. Select Command Console in the window.

5. Do one of the following:

- On Windows NT 4.0, click *Add/Remove*.
- On Windows 2000 and Windows Server 2003 (32-bit), click *Change/Remove*.

6. Follow the instructions on the screen.

**Note:** This procedure removes only the Command Console Client (SWCC Navigation Window). You can remove the HSG80 Client by using the Add/Remove program.

---

## Where to Find Additional Information

You can find additional information about SWCC by referring to the online Help and to *HP StorageWorks Command Console Version 2.5, User Guide*.

### About the User Guide

*HP StorageWorks Command Console Version 2.5, User Guide* contains additional information on how to use SWCC. Some of the topics in the user guide are the following:

- About AES
- Adding Devices
- Adding Virtual Disks
- Setting Up Pager Notification
- How to Integrate SWCC with Insight Manager
- Troubleshooting Information

### About the Online Help

Most of the information about the Client is provided in the online Help. Online Help is provided in two places:

- Navigation Window – Online Help provides information on pager notification and a tour of the Command Console Client, in addition to information on how to add a system to the Navigation Tree.
- Storage Window – Online Help provides detailed information about the Storage Window, such as how to create virtual disks.

# SWCC Agent in TruCluster Environment



This appendix describes how to set up SWCC Agent on TruClusters Tru64 UNIX Version 5.X

## SWCC Overview

The SWCC is a graphical user interface (GUI) for managing StorageWorks Redundant Array of Independent Disks (RAID) array products from a client (console) running on Microsoft Windows NT 4.0 with Service pack 4 or later, or Windows 2000. With SWCC you can configure virtual disks, receive notification of events, and monitor your storage systems.

The SWCC consists of a client and an agent. The client GUI runs on Windows NT 4.0 or Windows 2000 and provides an easy method of configuring, operating, monitoring, and troubleshooting your storage subsystem.

The agent is a companion program that is installed on one of the supported operating systems, in this case, Tru64 UNIX, to which a controller is connected. The agent lets the client communicate with the storage system over a TCP/IP network.

## Running the SWCC Agent under ASE Services (UNIX 4.0 only)

In a TruCluster environment, use ASE Services to monitor the state of an application. ASE Services can be used to automatically failover an application in the event that the node running the application fails. Create this “Service” by running the “`asemgr`” utility and answering the appropriate questions.

There are two steps to creating an ASE Service:

1. Create the Start/Stop scripts.
2. Create the service.

You can accomplish both steps by creating the *Start/Stop* scripts and then running the `asemgr` to finish creating the service.

---

**Note:** When the ASE Director is started (for example during boot up) ASE Services run the `Stop` scripts for all defined services by default. It then runs the `Start` scripts for all defined services. When writing `Start/Stop` scripts, one thing that should be apparent is scripts that toggle a program’s run state will exhibit erratic behavior.

---

## Creating the Start/Stop Script

To create Start/Stop scripts outside of the asemgr utility use your favorite editor. The Start script should start the program.

```
/usr/opt/SWCC520/scripts/steamstart
#!/bin/sh
#
# Simple start script for steamd.
#
PATH=.::/etc:/bin:/usr/bin:/usr/ucb:/sbin
export PATH
SD="/usr/opt/SWCC520/bin/steamd"
grep -qsw steam /etc/inittab
if [ $? -ne 0 ]
then
echo "steam:3:respawn:$SD sF" >> /etc/inittab
init q
else
exit 0
fi
```

An example of the Stop script follows which can be located in the /usr/opt/SWCC520/scripts directory. The keyword “steam” is searched for and if found the stream editor “sed” is invoked to remove it. The init daemon is then sent the “q” option to reload the inittab file. Since the “steam” line has been removed, initd will kill the steamd program.

```
/usr/opt/SWCC520/scripts/steamstop
#!/bin/sh
#
# Simple stop script for steamd.
#
PATH=.::/etc:/bin:/usr/bin:/usr/ucb:/sbin
export PATH
SD="/usr/opt/SWCC520/bin/steamd"
grep -qsw steam /etc/inittab
if [ $? -eq 0 ]
then
    sed '/steamd/d' /etc/inittab > /etc/inittab.tmp
    mv /etc/inittab.tmp /etc/inittab
    init q
else
    exit 0
fi
```

## Use the ASEMGR Utility to Create the Service

With the Start/Stop scripts created and run, complete the process by using the asemgr utility to create the service. A sample session follows:

```
tincp4>asemgr
```

### TruCluster Production Server (ASE)

```
Enter your choice: m
/ Select "m" to create a new service. /
ASE Main Menu
a) Managing the ASE -->
m) Managing ASE Services -->
s) Obtaining ASE Status -->
x) Exit ?) Help
```

### Managing ASE Services

```
Enter your choice [q]: c
/ Select "c" for a new service. /
Service Configuration
Enter your choice [q]: a
Adding a service
Enter your choice [1]: 3
```

When adding a service pick “3” even though the tendency would be to pick “2” a disk service which reserves a specific disk device. Since the Communication LUN may not always be the same, reserving a specific LUN would not be a good idea.

- c) Service Configuration -->
- r) Relocate a service
- on) Set a service on line
- off) Set a service off line
- res) Restart a service
- s) Display the status of a service
- a) Advanced Utilities -->
- q) Quit (back to the Main Menu)
- x) Exit ?) Help
- a) Add a new service
- m) Modify a service
- o) Modify a service without interrupting its availability
- d) Delete a service
- s) Display the status of a service
- c) Display the configuration of a service
- q) Quit (back to Managing ASE Services)
- x) Exit ?) Help

Select the type of service:

- 1) NFS service
- 2) Disk service
- 3) User-defined service
- 4) DRD service
- 5) Tape service
- q) Quit without adding a service
- x) Exit ?) Help

You are now adding a new user-defined service to ASE.

User-defined Service Name

The name of a user-defined service must be a unique service name within the ASE environment.

Enter your choice [x]: 1

As a minimum this is needed.

Enter your choice [x]: f

Modifying user-defined scripts for `steam`:

- 1) Start action
- 2) Stop action
- 3) Add action
- 4) Delete action
- 5) Check action
- x) Exit - done with changes

Modifying the start action script for `steam`:

- f) Replace the start action script
- e) Edit the start action script
- g) Modify the start action script arguments [steam]
- t) Modify the start action script timeout [60]
- r) Remove the start action script
- x) Exit - done with changes

Enter the full pathname of your start action script or “default” for the default script (x to exit):

/usr/opt/SWCC520/scripts/steamstart

This is the absolute path to the script illustrated previously.

Enter your choice [x]: x

Enter your choice [x]: 2

As a minimum, this is needed.

Modifying the start action script for `steam`:

- f) Replace the start action script
- e) Edit the start action script
- g) Modify the start action script arguments [steam]
- t) Modify the start action script timeout [60]
- r) Remove the start action script
- x) Exit - done with changes

Modifying user-defined scripts for `steam`:

- 1) Start action
- 2) Stop action
- 3) Add action
- 4) Delete action
- 5) Check action
- x) Exit - done with changes

Enter your choice [x]: f

Enter the full pathname of your stop action script or “default” for the default script (x to exit): /usr/opt/SWCC520/scripts/steamstop

This is the absolute path to the script illustrated previously.

```
Enter your choice [x]: x  
Enter your choice [x]: x
```

Modifying the stop action script for `steam`:

- f) Replace the stop action script
- e) Edit the stop action script
- g) Modify the stop action script arguments [steam]
- t) Modify the stop action script timeout [60]
- r) Remove the stop action script
- x) Exit - done with changes

Modifying the stop action script for `steam`:

- f) Replace the stop action script
- e) Edit the stop action script
- g) Modify the stop action script arguments [steam]
- t) Modify the stop action script timeout [60]
- r) Remove the stop action script
- x) Exit - done with changes

Modifying user-defined scripts for `steam`:

- 1) Start action
- 2) Stop action
- 3) Add action
- 4) Delete action
- 5) Check action
- x) Exit - done with changes

## Selecting an Automatic Service Placement (ASP) Policy

Enter your choice [b]: b

Selecting the Balanced option relocates steam to the least busy node should it be necessary. You may want to specify favored nodes. Do you want ASE to consider relocating this service to another member if one becomes available while this service is running (y/n/?): n

If "b" was NOT picked for the previous question, "y" should be picked for this one.

Enter 'y' to add Service 'steam' (y/n): y

Adding service...

Starting service...

Service steam successfully added...

Enter your choice [q]: q

Select the policy you want ASE to use when choosing a member to run this service:

- b) Balanced Service Distribution
- f) Favor Members
- r) Restrict to Favored Members
- ?) Help

### Service Configuration

- a) Add a new service
- m) Modify a service
- o) Modify a service without interrupting its availability
- d) Delete a service
- s) Display the status of a service
- c) Display the configuration of a service
- q) Quit (back to Managing ASE Services)
- x) Exit ?) Help

### Managing ASE Services

Enter your choice [q]: q

TruCluster Production Server (ASE)

Enter your choice: x

tincp4>

If everything was entered correctly the steamd should be running on one of the nodes in the cluster. If that node should go down, the steamd will automatically be relocated to another node in the cluster.

---

**Note:** If the SWCC Agent is running on a node which has a LUN reserved to it, and that Agent fails over to another node, you will no longer be able to monitor that LUN. If a disk service is relocated, run `/usr/opt/SWCC520/scripts/swcc_config` and select Option 2 to reconfigure your storage. Then restart the Agent to have these changes take effect.

---

- c) Service Configuration -->
- r) Relocate a service
- on) Set a service on line
- off) Set a service off line
- res) Restart a service
- s) Display the status of a service
- a) Advanced Utilities -->
- q) Quit (back to the Main Menu)
- x) Exit ?) Help

#### ASE Main Menu

- a) Managing the ASE -->
- m) Managing ASE Services -->
- s) Obtaining ASE Status -->
- x) Exit ?) Help

## Running the SWCC Agent on a V5.x Cluster

This section is intended for users of Tru64 UNIX V5.x. Running SWCC on a V5.x Cluster has the following requirements:

**Table 20: Configuration Requirements**

Requirement	Description
Operating System	Tru64 UNIX Version 5.0 or higher and TruCluster Server Version 5.0A or higher.
Supported Controllers and Minimum Firmware Versions	<p>The supported firmware has locking features that prevent data corruption resulting from access by multiple systems. The supported combinations of controllers and minimum firmware versions are as follows.</p> <p>HSZ40 3.7; HSZ50 5.7; HSZ70 7.7          HSZ80 8.3-1; HSG80 8.5; HSG60 8.5L          Earlier versions of the firmware are explicitly not supported.</p> <p>HSZ40 3.7          HSZ50 5.7          HSZ70 7.7          HSZ80 8.3-1          HSG80 8.5          HSG60 8.5L          Earlier versions of the firmware are explicitly not supported.</p>
Disk Space Required for Agent	2 MB free space in /tmp directory. 4 MB free space for the Agent installation directory.
Skill Level Requirements	This procedure is intended for experienced system administrators or programmers. Familiarity with shell programming and access to the root of the TruCluster Server Cluster is required.

The SWCC Tru64 UNIX agent installation script is not cluster-aware. It adds swcc and S500swcc entries to the /sbin/init.d and /sbin/rc3.d directories, respectively, and starts the steamd daemon on the system.

In addition, the SWCC configuration script (/usr/opt/SWCCx/scripts/swcc\_config) and initialization file (/usr/opt/SWCCx/etc/storage.ini) are not context-dependent symbolic links (CDSLs) and are shared by all cluster members. Therefore,

installing and configuring the SWCC agent on one cluster member configures the agent using the storage information for that cluster member, but the configuration information is then used on all cluster members.

It can be confusing if you run the SWCC agent on multiple cluster members and those members do not have the same view of the storage. For example, assume that you have used the selective storage presentation (SSP) feature (ENABLE\_ACCESS\_PATH=, DISABLE\_ACCESS\_PATH=) to control which hosts have access to disk units D8, D9, and D10. Further assume that cluster member swiss can access these units, and that cluster member rye cannot. When you use the hwmgr command on each member to look at the available devices, member swiss can see units D8, D9, and D10, indicated by the user-assignable identifier (IDENTIFIER=x), but member rye cannot:

```
# hostname
swiss.zk3.dec.com
# hwmgr -view devices
HWID: Device Name Mfg Model Location
-----
54: /dev/disk/dsk0c DEC HSG80 IDENTIFIER=1
55: /dev/disk/dsk1c DEC HSG80 IDENTIFIER=2
56: /dev/disk/dsk2c DEC HSG80 IDENTIFIER=3
57: /dev/disk/dsk3c DEC HSG80 IDENTIFIER=4
58: /dev/disk/dsk4c DEC HSG80 IDENTIFIER=5
59: /dev/disk/dsk5c DEC HSG80 IDENTIFIER=6
60: /dev/disk/dsk6c DEC HSG80 IDENTIFIER=7
62: /dev/cport/scp0 HSG80CCL bus-4-targ-0-lun-0
68: /dev/kevm
108: /dev/disk/floppy1c 3.5in floppy fdi0-unit-0
119: /dev/disk/cdrom1c DEC RRD47 (C) DEC bus-0-targ-5-lun-0
120: /dev/disk/dsk7c COMPAQ BB00911CA0 bus-1-targ-0-lun-0
125: /dev/dmapi/dmapi
126: /dev/disk/dsk8c DEC HSG80 IDENTIFIER=8
127: /dev/disk/dsk9c DEC HSG80 IDENTIFIER=9
128: /dev/disk/dsk10c DEC HSG80 IDENTIFIER=10
```

```
# hostname
rye.zk3.dec.com
# hwmgr -view devices
HWID: Device Name Mfg Model Location
-----
4: /dev/kevm
44: /dev/disk/floppy0c 3.5in floppy fdi0-unit-0
53: /dev/disk/cdrom0c DEC RRD47 (C) DEC bus-0-targ-5-lun-0
54: /dev/disk/dsk0c DEC HSG80 IDENTIFIER=1
55: /dev/disk/dsk1c DEC HSG80 IDENTIFIER=2
56: /dev/disk/dsk2c DEC HSG80 IDENTIFIER=3
57: /dev/disk/dsk3c DEC HSG80 IDENTIFIER=4
58: /dev/disk/dsk4c DEC HSG80 IDENTIFIER=5
59: /dev/disk/dsk5c DEC HSG80 IDENTIFIER=6
60: /dev/disk/dsk6c DEC HSG80 IDENTIFIER=7
61: /dev/ntape/tape0 DEC TLZ10 (C) DEC bus-0-targ-4-lun-0
62: /dev/cport/scp0 HSG80CCL bus-4-targ-0-lun-0
125: /dev/dmapi/dmapi
```

However, because SWCC looks at storage from the controller's viewpoint, units D8, D9, and D10 appear in the Storage Window for rye even though member rye cannot directly access these units, as shown in the Example of Confusing Storage Display.

As another example, if you use the

`/usr/opt/SWCCx/scripts/swcc_config` configuration tool to make a change to the storage subsystem information and restart the agent on one member, you must also use the `/usr/opt/SWCCx/scripts/swcc_config` configuration tool to restart the agent on all other members. If you do not do this, the Navigation Window for the two members can show a discrepancy.

## Problems with Running the Agent on Multiple Clusters

Running the agent on multiple cluster members can cause problems in operational ease-of-use. For example, if you make changes with the agent on one cluster member, those changes are not immediately reflected by the agents on other cluster members, so what you see in the Navigation Window might not be what you expect. Running the agent on multiple cluster members is not harmful, but it can be confusing, depending on your configuration.

The required ACS firmware version for the controller implements a locking mechanism that prohibits data corruption caused by simultaneous system access, either from cluster members or standalone systems. Your data integrity is protected.

The following is recommended for running SWCC in a TruCluster Server environment. The controller recommendations are specific to the HSG80, HSG60, and HSZ80 controllers:

- Configure your hardware as described in the TruCluster Server Hardware Configuration manual.
- Use multiple-bus failover mode so that all units are available to both ports of an HSG80/HSG60/HSZ80 controller.
- Set the HSG80/HSG60 unit offset to 0 for all connections.
- Set the HSG80/HSG60/HSZ80 access path so that the appropriate cluster member connections are enabled.
- Make sure that the HSG80/HSG60/HSZ80 Command Console LUN (CCL) is enabled. We recommend that you use the controller in SCSI-3 mode so that the CCL is always located at LUN0.
- If you use selective storage presentation (SSP) on the HSG80/HSG60 controller, or zoning on a Fibre Channel switch, remember that this selective view of your storage is reflected in the Navigation Window.
- Install the SWCC agent on one cluster member. (The SWCC installation script does not allow you to install it on more than one member at a time.)
- Edit the `/sbin/init.d/swcc` startup script to probe for a TruCluster Server environment.
- Use the cluster application availability (CAA) scripts described in this procedure to start and stop the SWCC agent, and to provide failover for the `swcc` resource.

- Configure the SWCC client to use a cluster alias as the address for the SWCC agent system to avoid having to change the client if the CAA swcc resource fails over to another member. You can use the default cluster alias, or any alias to which all of the cluster members belong.

## Configure the Controller

The Hardware Configuration manual describes how to configure the hardware in a TruCluster Server environment. This manual includes complete information on setting up a controller in a shared SCSI or Fibre Channel environment. If applicable to your environment, complete the tasks described in the Using Fibre Channel Storage chapter. That chapter explains the use of Fibre Channel switches and hubs, zoning switches and selective storage presentation (SSP). It also describes example Fibre Channel configurations supported by TruCluster Server.

Use the Hardware Configuration manual to configure your hardware before you implement this appendix. Although you perform some hardware configuration in this appendix, do so after your initial hardware configuration is complete.

## Use Multiple-Bus Failover Mode

In multiple-bus failover mode, all unit numbers (0 through 199) are potentially visible on all four controller ports in a dual-redundant configuration. With multiple-bus failover:

- The host controls the failover by accessing units over a different path or causing the access to the unit to be through the other controller.
- An active controller causes a failover to the other controller if the controller recognizes the loss of the switch, hub, or cable to a controller port.
- Each cluster member system can have two or more (fabric only) KGPSA host bus adapters (multiple paths to the storage units).

You set multiple-bus failover from the controller maintenance terminal with the set multibus copy = thisother command. Ensure that you copy the configuration information from the controller that you know has a good array configuration.

---

**Note:** When you enter the command to set multiple-bus failover and copy the configuration information to the other controller, the other controller will restart. The restart may set off the audible alarm (which is silenced by pressing the button on the environmental monitoring unit (EMU)). The command line interpreter (CLI) will display an event report, and continue reporting the condition until you clear the condition with the `clear cli` command.

---

```
HSG80> set multibus copy = this
HSG80> show this
Controller:
HSG80 ZG02700662 Software V85F-0, Hardware E05
NODE_ID = 5000-1FE1-0006-3F10
ALLOCATION_CLASS = 0
SCSI_VERSION = SCSI-3
Configured for MULTIBUS_FAILOVER with ZG02700619
In dual-redundant configuration
Device Port SCSI address 7
Time: 01-MAR-2001 17:48:32
Command Console LUN is lun 0 (NOIDENTIFIER)
```

## Verify That the HSG80/HSG60 Unit Offsets Are Zero

In multiple-bus failover mode, the default offset is 0 for all host connections. However, if a controller pair is switched from transparent failover mode to multiple-bus failover mode, the unit offsets for transparent mode remain in effect, so a connection can have a non-zero unit offset. For each connection to your cluster, verify that the unit offset is 0. If the offsets are incorrect, set them and then restart both controllers.

```
HSG80> show connection
```

### Connection

Name	Operating system	Controller	Port	Address	Status	Offset
RYE_1	TRU64_UNIX	OTHER	2	011400	OL	other 0
						HOST_ID=2000-0000-C922-4AAC ADAPTER_ID=1000-0000-C922-4AAC
RYE_2	TRU64_UNIX	THIS	1	011400	OL	this 0
						HOST_ID=2000-0000-C922-4AAC ADAPTER_ID=1000-0000-C922-4AAC
RYE_3	TRU64_UNIX	OTHER	1	011400	OL	other 0
						HOST_ID=2000-0000-C922-4AAC ADAPTER_ID=1000-0000-C922-4AAC
RYE_4	TRU64_UNIX	THIS	2	011400	OL	this 0
						HOST_ID=2000-0000-C922-4AAC ADAPTER_ID=1000-0000-C922-4AAC
SWISS_1	TRU64_UNIX	OTHER	2	011200	OL	other 0
						HOST_ID=2000-0000-C924-4B7B ADAPTER_ID=1000-0000-C924-4B7B
SWISS_2	TRU64_UNIX	THIS	1	011200	OL	this 0
						HOST_ID=2000-0000-C924-4B7B ADAPTER_ID=1000-0000-C924-4B7B
SWISS_3	TRU64_UNIX	OTHER	1	011200	OL	other 0
						HOST_ID=2000-0000-C924-4B7B ADAPTER_ID=1000-0000-C924-4B7B
SWISS_4	TRU64_UNIX	THIS	2	011200	OL	this 0
						HOST_ID=2000-0000-C924-4B7B ADAPTER_ID=1000-0000-C924-4B7B

If an offset is not zero, you can change it as follows:

```
HSG80> set RYE_1 unit_offset = 0
HSG80> restart other
HSG80> restart this
```

## Use SCSI-3 Mode

A logical unit number (LUN) is an address of a logical unit on a virtual disk. The Command Console LUN (CCL), also called a communications LUN, is a special logical unit number that is used to communicate with the controller to set up the SWCC. We recommend that you use the controller's SCSI-3 mode, which forces the CCL to be always enabled, at LUN 0, and does not allow it to float.

Use the `show this` and `show other` commands to determine which SCSI mode is being used.

```
HSG80> show this
```

Controller:

```
HSG80 ZG02700662 Software V85F-0, Hardware E05
NODE_ID = 5000-1FE1-0006-3F10
ALLOCATION_CLASS = 0
SCSI_VERSION = SCSI-3
Configured for MULTIBUS_FAILOVER with ZG02700619
In dual-redundant configuration
Device Port SCSI address 7
Time: 01-MAR-2001 17:48:32
Command Console LUN is lun 0 (NOIDENTIFIER)
:
:
HSG80> show other
```

Controller:

```
HSG80 ZG02700619 Software V85F-0, Hardware E05
NODE_ID = 5000-1FE1-0006-3F10
ALLOCATION_CLASS = 0
SCSI_VERSION = SCSI-3
Configured for MULTIBUS_FAILOVER with ZG02700662
In dual-redundant configuration
Device Port SCSI address 6
Time: 02-MAR-2001 18:07:48
Command Console LUN is lun 0 (NOIDENTIFIER)
```

## Install and Run the Agent on One Cluster Member

We recommend that you run the SWCC agent on only one cluster member and use CAA for failover. The SWCC Navigation Window treats cluster members as standalone systems. That is, the Navigation Window does not include a cluster management object; you initiate all management tasks at the system (member) level. If you were to add each cluster member to the Navigation Window, each member would be displayed as a standalone system and would appear to have ownership of the controllers attached to it, which might be confusing at first glance.

Follow these steps to install the agent software on a TruCluster Server cluster member:

1. Insert the storage subsystem CD-ROM in the CD-ROM drive.
2. Mount the CD-ROM with the command:

```
# mount -r -t cdfs /dev/disk/cdrom0c /mnt
```

where *cdrom0c* is the name of the CD-ROM drive.

3. Change to the Command Console directory with the following command:

```
# cd /mnt/agent
```

4. Execute the installation script with this command:

```
# setld -1
```

5. Follow the online instructions to install the software.

The agent installation procedure requires details about the configuration, including the information listed in the following table:

**Table 21: Required Configuration Details**

Property	Description
The host name of the client systems	Network names for the computers (PCs) on which the Command Console Client (the GUI) will be running.
Password	Required of the client system to gain configuration access.
Access level	The amount of control to be granted to the client system. If you want to restrict the type of access from a client, you must specify a password. The possible options are: 0 = Overall status 1 = Detailed status 2 = Configuration and Status

**Table 21: Required Configuration Details (Continued)**

Property	Description
Notification schemes	<p>The mechanism by which to notify you of a change in status. The possible options are:</p> <p>0 = No notification over the network. Local notification through e-mail and an entry in the system error log file.</p> <p>1 = Notification sent to the client via TCP/IP. Local notification sent through e-mail and an entry in the system error log file.</p> <p>2 = Notification sent to the Simple Network Management Protocol (SNMP) monitor. Local notification sent through e-mail and an entry in the system error log file.</p> <p>3 = Notification sent to client via TCP/IP and the SNMP monitor. Local notification sent through e-mail and an entry in the system error log file.</p>

**Table 21: Required Configuration Details (Continued)**

Property	Description
Subsystem name	A unique, user-defined name for each controller in a single-controller storage configuration, and for each controller pair in a dual-redundant configuration. Try to use a descriptive name that will clearly identify the controller or its location.
E-mail destination for notification	<p>An e-mail account that will receive error and status information. Specifying an e-mail account is optional, and requires that mail be already configured on the TruCluster Server cluster. See the Cluster Administration manual for information on configuring mail.</p> <p>The agent sends a mail notification in a form similar to the following:</p> <p>From: StorageWorks_Agent@deli.zk3.dec.com    Date: Tue, 6 Mar 2001 03:29:13 -0500 (EST)    Sub: StorageWorks Agent on swiss.zk3.dec.com    This is a message from your StorageWorks    RAID Manager Software. A RAID system status    change has been detected. Corrective    action may be required! Refer to the    following status information and your Agent    Log file for more direction.    Hostname: swiss.zk3.dec.com    CRITICAL: Unable to open device -    scp0 (SP_MONITOR: MonitorSubsys)</p>
Cluster member status	<p>If you are installing this agent in a TruCluster Server environment, answer yes.</p> <p>This option invokes the steamd -s option, which forces the steamd daemon to rescan the SCSI busses on startup to detect any change in device configuration.</p>

## Example of Installing the Agent on a Cluster Member

The example of Installing the Agent, using Version SWCC5xx shows a complete log of the agent installation on a TruCluster Server cluster member.

### Example of Installing the Agent, Using Version SWCC5xxx

```
# mount -r -t cdfs /dev/disk/cdrom0c /mnt
# cd /mnt/agent
# ls
SWCC5xx instctrl
# setld -l .
*** Enter subset selections ***
```

The following subsets are mandatory and will be installed automatically unless you choose to exit without installing any subsets:

```
* SWCC Agent 2.3 For Tru64 Unix
```

You may choose one of the following options:

- 1) ALL of the above
- 2) CANCEL selections and redisplay menus
- 3) EXIT without installing any subsets

```
Estimated free diskspace(MB) in root:323.9 usr:1285.2 var:1756.5
```

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6): 1

You are installing the following mandatory subsets:

```
SWCC Agent 2.3 For Tru64 Unix
```

You are installing the following optional subsets:

```
Estimated free diskspace(MB) in root:323.9 usr:1285.2 var:1756.5
```

Is this correct? (y/n): y

```
setld:
```

```
Checking file system space required to install selected subsets:  
File system space checked OK.  
1 subsets will be installed.  
Loading subset 1 of 1 ...  
SWCC Agent 2.3 For Tru64 Unix  
Copying from . (disk)  
Verifying  
1 of 1 subsets installed successfully.  
Setting up daemon for V5.x.  
NOTE: 'catman' is currently running, so I can't update your  
'whatis' database. After this installation run 'catman -w' if  
you want to update 'whatis'.  
----- CLIENT.INI -----  
The 'client.ini' file stores information about PC's running the  
SWCC Graphical User Interface (GUI). To allow client systems  
running the GUI access to your StorageWorks RAID subsystem you  
will need to enter the following information for each client.  
* Clients hostname.  
* Clients access privilege.  
* Clients notification level.  
For more information see man page steamd.8.  
Press 'Return' to Continue  
Adding a Client system  
-----  
Enter the host name of the Client system:adminpc.zk3.dec.com  
Enter the subsystem access privilege (which controls the level  
of access  
a client has for communicating with the storage subsystems).  
The possible options are:  
0 = Overall Status  
1 = Detailed Status  
2 = Configuration + Status  
Enter Access Level (0, 1, 2): 2
```

The Agent server can notify a client when an error condition occurs.

Notification schemes available are:

0 = No Error Notification  
1 = Notification via a TCP/IP Socket  
2 = Notification via the SNMP protocol  
3 = Notification via both TCP/IP and SNMP

Enter Error Notification Level (0, 1, 2, 3): 1

Review Client Information--

name: adminpc.zk3.dec.com

access level: 2

error notification: 1

Is this information correct? [y,n] y

Would you like to add another client? [y,n] n

----- PASSWORD -----

In order for a user of the SWCC client GUI to do remote configurations, a password is required. If this installation is superseding an older version and the old password is found; you will be asked if you want to use it. If you don't use it you will be asked for a new password For more information see man page steamd.8.

Press 'Return' to Continue

Changing Client Access Password

-----

Enter Password (4-16 Chars)

OR

Press 'ESC' RETURN to Quit:

Re-Enter Password:

PASSWORD CHANGED!

Press RETURN to continue:

----- STORAGE.INI -----

The 'storage.ini' file stores information about RAID devices connected to your server. The SWCC agent reads this file at startup. If this installation supersedes an older version and the storage.ini file is found; you will be given the option to use it. If not; the following information will be needed for each RAID subsystem you want the agent to interact with.

- \* RAID subsystem name.
- \* Monitoring Interval in seconds.

For more information see man page steamd.8.

Press 'Return' to Continue

Scanning your system. This may take awhile depending on the size of your system. Please Stand-By.

\* \* \*

Adding a Storage Subsystem

-----

Enter a name for the FIRST HSG80 I found on this system:  
delihsg80

The Agent monitors RAID systems at regular intervals. You specify this polling interval in seconds ( 0 = No monitor; 300 = every 5 minutes. )

Enter Monitoring Interval for delihsg80: 300

Review Storage Subsystem Information--

name: delihsg80

Monitor Interval: 300

Is this information correct? [y,n] y

```
----- NOTIFY.INI -----
```

The 'notify.ini' file stores mail notification information. If you and/or anyone else should be notified via Email when a RAID subsystem event occurs, enter their Email address in the next screen. If this installation supersedes an older version and an existing notify.ini is found; you will be given the option to use it. If not; the following information will be needed for each person you want notified.

```
* Email address of person to notify.  
* Notification Level.
```

For more information see man page steamd.8.

Press 'Return' to Continue

Change State of Email Notification

```
-----  
Email Notification is currently DISABLED. Would you  
like to enable it?
```

Enter [y/n] :y

Email Notification is now ENABLED.

Adding a Storage Email Notification User

```
-----  
Enter the Email address of a person to notify : sysadm@zk3.dec.com
```

The Agent has 3 levels of mail notification ranging from Critical

to Critical and Warning to Critical and Warning and Informational.

1 = Fatal Errors

2 = Warnings and Fatal Errors

3 = Information, Warnings and Fatal Errors

Enter level of notification : 1

--Review Email Notification Information--

Email: sysadm@zk3.dec.com

Level: Critical Errors Only (1)

Is this information correct? [y,n] y

Would you like to add another notify user? [y,n] n

----- toggle -----

The SWCC Agent can run on a stand alone server or in a Tru64 TruCluster environment. The next question you will need to answer determines arguments supplied to the Agent at startup. If you are installing this Agent in a TruCluster environment answer 'Y' for yes. If you are not running TruCluster answer 'N' for no.

REMEMBER: Only one instance of the Agent should run in a cluster unless all HSZs and HSGs in the cluster are running firmware versions that support the SWCC lock bit. See your HSx documentation for more information.

For more information on switches see man page steamd.8.

Press 'Return' to Continue

Are you running the Agent in a cluster? [y,n] y

Setting 's' switch

SWCC5xx software installed successfully.

steamd config: Device Failover disabled

steamd config: Bus Scanning enabled, forcing initial scan

Agent Version: 2.3 (Build 77)

GetStgParms>> Scanning the bus for devices

Configuring "SWCC Agent 2.3 For Tru64 Unix" (SWCC5xx) on member0

Configuring "SWCC Agent 2.3 For Tru64 Unix" (SWCC5xx) on member1

Configuring "SWCC Agent 2.3 For Tru64 Unix" (SWCC5xx) on member2

# cd

# umount /mnt

# ps agx | grep steamd

540567 pts/1 S 0:00.11 /usr/opt/SWCC5xx/bin/steamd sf

540850 pts/1 S + 0:00.00 grep steamd

## Create the CAA Action Script

You can have CAA control the steamd daemon and fail it over to another member when necessary.

The CAA action script invokes the startup and shutdown functions in the shell script. If the cluster member on which the steamd daemon is running fails or is shut down, CAA relocates the daemon entry to another member. The action script is called swcc.scr here but you can name it after the CAA resource name of your choice.

Save the action script shown in Complete CAA Action Script to the file /var/cluster/caa/script/swcc.scr, or another filename of your choice. (We suggest that you use this action script as-is until you get the swcc resource working.) This file must exist for the CAA commands to complete successfully. Make sure that this file is executable, because it is called by CAA.

The example script assumes that you are using the SWCC5xx version of the agent. If you have a later version of the agent, as identified in the /usr/opt directory, update the following line in the script.

```
APPDIR="/usr/opt/SWCC5xx" # Application home directory
```

The example directs output to log files, using the directory /var/cluster/caa/log. These log files can be useful for debugging purposes. Change the directory as appropriate.

## Complete CAA Action Script

```
#!/usr/bin/ksh -p
#
svcName="swcc"># Servicename
CAA_ADMIN="root"# Account to receive CAA mail
CAALOGDIR="/var/cluster/caa/log"# Directory for logfiles
ACTION=$1# Action (either start or stop)
LOG="${CAALOGDIR}/${ACTION}_${svcName}.$$$" # Destination for
script output
#LOG="/dev/console"
#
PROBE_PROCS="steamd"# Processes to probe
START_APPCMD="startswcc"# Application startup cmd
STOP_APPCMD="zapdaemon steamd"# Application stop cmd
APPDIR="/usr/opt/SWCC5xx"# Application home directory
ADVFSDIRS=" "# Application directories to
#
# Custom procedures for SWCC
#
startswcc () {
if [ -f ${APPDIR}/bin/steamd ]
then
nohup ${APPDIR}/bin/steamd SF >> ${LOG} &
else
echo "steamd executable missing, aborting startup..." >> ${LOG}
return 2
fi
}
# Application specific stuff
#
#
FUSER="/usr/sbin/fuser" # Command to use for closing
EVMPOST="/usr/bin/evmpost -p 650 -a"# EVM command to post events
#
```

```
export START_APPCMD STOP_APPCMD APPDIR ADVFSDIRS PROBE_PROCS
#
# 8<-----8<----- End Custom variables
8<-----8<-----
#
# Static variables
#
PATH=/sbin:/usr/sbin:/usr/bin:/usr/lbin
TERM=vt100
SHELL=/usr/bin/sh
HOME=/
USER=root
LOGNAME=root
HOST=`/bin/hostname`
umask 117
cd ${DIR}
OLDPWD=`pwd`
export ACTION DIR PATH TERM SHELL HOME USER LOGNAME HOST OLDPWD
#
# Frequently used procedures
#
checkdaemon () {
R=`ps -o command -A | grep $1 | grep -v grep | wc -l`
return $R
}
#
killdaemon () {
kill -TERM `ps -o pid,command -A | grep $1 | grep -v grep | awk
'{print $1}'` 
while checkdaemon $1; [ $? -gt 0 ]; do
sleep 1
done
}
#
```

```
# Kill (-9) a given process using brutal force
#
zapdaemon () {
for i in ${1}
do
kill -9 `ps -o pid,command -A | grep ${i} | grep -v grep | awk
'{print $1}'` \
checkdaemon ${i}
if [ $? -ne 0 ]; then
echo "Retrying to kill process ${i} "
kill -9 `ps -o pid,command -A | grep ${i} | \
grep -v grep | awk '{print $1}'` \
checkdaemon ${i}
if [ $? -ne 0 ]; then
echo "Process ${i} (PID: `ps -o pid,command -A | grep ${i} | \
grep -v \
grep | awk '{print $1}'` ) \
seems to be stubborn, please investigate"
exit 2
fi
fi
done
}
#
# If failure, post EVM event
#
postevent () {
${EVMPPOST} "*URGENT* CAA Script Failure, Cannot ${2} ${1} .
Aborting ${2} "
echo "Cannot ${2} ${1} . Aborting ${2} at `date`." >> ${LOG}
echo "Cannot ${2} ${1} . Aborting ${2} at `date`." | \
/usr/bin/mailx -s "*URGENT* CAA Script Failure" $CAA_ADMIN
return 2
}
#
```

```
# Probe for a running process/application
#
probeapp () {
ps -o command -A | grep $1 | grep -v grep > /dev/null 2>&1
if [ $? -ne 0 ]; then
echo "Cannot probe process ${1} . Posting EVM event."
${EVMPPOST} "*URGENT* CAA Script Failure, Probing of process ${1}
failed"
exit 2
fi
echo "Application ${1} is alive"
}
#
# Main section
#
# Start section
#
case $1 in
'start')
echo "" >> ${LOG}
echo "" "Start action script for service : ${svcName} \
`/bin/date +"%A %d %B %H:%M:%S"`` "" >> ${LOG}
#
# Start SWCC
#
echo "Starting SWCC . . ." >> ${LOG}
cd $APPPDIR
$START_APPCMD>> ${LOG}
if [ $? -ne 0 ]; then
postevent "SWCC startup" start
exit 2
fi
echo "Started SWCC steamd">> ${LOG}
#
```

```
# All done ...
#
${EVMPPOST} "Start action script for service ${svcName} DONE"
echo "" "Start action script for service ${svcName} DONE, \
`/bin/date +"%A %d %B %H:%M:%S"` "" >> ${LOG}
echo "" >> ${LOG}
exit 0
#
;;
#
# Stop section
#
'stop')
echo "" >> ${LOG}
echo "" "Stop action script for service : ${svcName} \
`/bin/date +"%A %d %B %H:%M:%S"` "" >> ${LOG}
#
# Stop SWCC
#
echo "Stopping SWCC ...">>> ${LOG}
$STOP_APPCMD>> ${LOG}
if [ $? -ne 0 ]; then
postevent "SWCC steamd" stop
exit 2
fi
$STOP_APPCMD2 >> ${LOG}
if [ $? -ne 0 ]; then
postevent "SWCC steamd" stop
exit 2
fi
echo "SWCC steamd shutdown done .">>> ${LOG}
${EVMPPOST} "Stop action script for service ${svcName} DONE"
echo "" "Stop action script for service ${svcName} DONE, \
`/bin/date +"%A %d %B %H:%M:%S"` "" >> ${LOG}
```

```
echo "" >> ${LOG}
exit 0
;;
#
# Probe if application is still alive
#
'check')
echo ""Probing SWCC daemons at \
`/bin/date +"%A %d %B %H:%M:%S"``" >> ${LOG}
for i in ${PROBE_PROCS}
do
probeapp ${i} >> ${LOG}
done
echo ""Probing SWCC daemons DONE at \
`/bin/date +"%A %d %B %H:%M:%S"``" >> ${LOG}
exit 0
;;
*) echo "usage: $0 {start|stop|check}"
exit 1
;;
esac
```

## Create the CAA Resource

Use the `caa_profile` command to create a CAA resource profile, and use the `caa_register` command to register the resource after you have created it.

This procedure uses `swcc` as the resource name for the purpose of the example. Use `application` as the resource type and specify the location of the action script.

If you want to start the resource automatically, regardless of whether it had been stopped or running before the reboot, set `auto_start` to `auto_start=1`. Set `auto_start=0` if you want to start the resource only if it had been running before the reboot.

An example `caa_profile -create` command is as follows:

```
# caa_profile -create swcc -a swcc.scr -t application -o as=1
```

You can specify either a full pathname for the script file, or its filename, in which case the `caa_profile` command looks for the file in the `/var/cluster/caa/script` directory.

By default, `caa_profile -create` enables `steamd` to run on all members with a placement policy of balanced. The `steamd` daemon can therefore run on any member of the cluster. If this default does not meet your needs, use the `-c` option to specify the members that can host the application resource and use the `-p` option to specify a placement policy.

The default check interval is 60 seconds. The check interval is the maximum amount of time that an application can be unavailable to clients before CAA attempts to restart it. Under some circumstances, this default can lead to a situation where a resource fails and is not restarted quickly enough. You need to determine whether the 60-second check interval is suitable.

If the caa\_profile -create command completes successfully, use the caa\_profile -print resource\_name command to verify the profile is as you intended:

```
# caa_profile -print swcc
NAME=swcc
TYPE=application
ACTION_SCRIPT=swcc.scr
ACTIVE_PLACEMENT=0
AUTO_START=1
CHECK_INTERVAL=60
DESCRIPTION=swcc
FAILOVER_DELAY=0
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=
OPTIONAL_RESOURCES=
PLACEMENT=balanced
REQUIRED_RESOURCES=
RESTART_ATTEMPTS=1
SCRIPT_TIMEOUT=60
```

## Register the CAA Resource

Use the caa\_register resource\_name command to register the resource with CAA:

```
# caa_register swcc
```

### Example of Creating, Registering, and Starting the CAA Resource

```
# caa_profile -create swcc -a swcc.scr -t application -o as=1
# caa_profile swcc -print
NAME=swcc
TYPE=application
ACTION_SCRIPT=swcc.scr
ACTIVE_PLACEMENT=0
AUTO_START=1
CHECK_INTERVAL=60
DESCRIPTION=swcc
FAILOVER_DELAY=0
FAILURE_INTERVAL=0
FAILURE_THRESHOLD=0
HOSTING_MEMBERS=
OPTIONAL_RESOURCES=
PLACEMENT=balanced
REQUIRED_RESOURCES=
RESTART_ATTEMPTS=1
SCRIPT_TIMEOUT=60
# caa_register swcc
# caa_stat swcc
NAME=swcc
TYPE=application
TARGET=OFFLINE
STATE=OFFLINE
# caa_start swcc
Attempting to start `swcc` on member `swiss`
Start of `swcc` on member `swiss` succeeded.
# caa_stat swcc
NAME=swcc
TYPE=application
TARGET=ONLINE
STATE=ONLINE on swiss
```

See caa\_profile(8) and caa\_register(8) for additional information.

## Viewing Events Posted by the Action Script

The action script shown in Complete CAA Action Script posts events to indicate state transitions and failures. You can use the event viewer to monitor these events. You can launch the event viewer through SysMan Menu or SysMan Station. See sysman(8) for additional information.

The following example shows sample event monitor output generated by the action script.

```
-----+
-----+
EventId Pri Timestamp Host Summary
-----+
234 650 14-Mar-2001 09:54:29 swiss EVM admin msg: Start action
script for service swcc DONE
231 650 14-Mar-2001 09:54:26 swiss EVM admin msg: Stop action
script for service swcc DONE
230 600 14-Mar-2001 09:54:23 swiss CAAD[1062599]: `swcc` on
`swiss` went OFFLINE unexpectedly
228 650 14-Mar-2001 09:54:23 swiss EVM admin msg: *URGENT* CAA
Script Failure, Probing of process steamd failed
207 650 14-Mar-2001 09:53:21 swiss EVM admin msg: Start action
script for service swcc DONE
```

## Edit the Startup Script

After you are satisfied that swcc is operating correctly under CAA control, edit the beginning of the /sbin/init.d/swcc file to test whether the swcc resource is registered with CAA, and exit if it is. (If you did not name the resource swcc when you created it, use the name you chose.)

```
#!/sbin/sh
#
/usr/bin/caa_stat -a swcc -g
is_registered=$?
if [ "$is_registered" = 0 ]
then
exit 0
fi
trap "echo" 2
IBASE="/usr/opt/SWCC5xx"
PATH=/sbin:/usr/sbin:/usr/bin:/bin
export PATH
SWITCH=`cat $IBASE/etc/drs.conf`
```

:

:

## Install and Configure the SWCC Client

The Command Console Client runs only on Microsoft Windows NT 4.0, service pack 4 or later, or Windows 2000. The client installation program is included on the Solution Software CD-ROM that is packaged in the Tru64 UNIX Platform Solution Kit.

---

**Note:** If a previous version of the Command Console Client is already installed on the PC, you must uninstall it before beginning the installation.

---

To install the client:

1. Insert the Solution SW CD-ROM from the platform kit into the CD-ROM drive of the PC that will run the client.

2. If the program does not start automatically, run the setup file D:\SWCC\CLIENT\INTEL\SETUP.EXE, where D: is the location of your CD-ROM drive. The installation program is self-extracting and by default stores the client software in the C:\Program Files\Compaq\SWCC directory.
3. You are presented with a list of available client applications. Choose the applications appropriate for your environment. For an HSG80/HSG60 environment, choose only the following:
  - SWCC Navigation Window
  - HSG80 Controller for ACS85 or HSG60 Controller
  - Optionally, Command Line Interface for HS\*

Be sure to choose HSG80 Controller for ACS87 and not HSG80 Controller.
4. Follow the installation script to complete the procedure.
5. If you are prompted to restart the system when the installation is complete, click Yes.
6. From the Windows NT or Windows 2000 Start menu, Command Console Client provides you with choices. To begin, click on StorageWorks Command Console. StorageWorks Command Console contains a Navigation Window that lets you connect to, monitor, and configure many subsystems.
7. From the File menu, select Add System. The Example Add System Window dialog box appears.
8. Enter the name of a cluster alias and click on Apply. You can use the default cluster alias, or any alias to which all of the cluster members belong.
9. Expand the system object for the cluster (remember that SWCC treats this as system, not as a cluster). Note that any storage systems detected by the SWCC agent running on the cluster are displayed.

## **Do Not Use `swcc_config` to Start, Stop, or Restart**

After you place the steamd daemon under CAA control and edit the /sbin/init.d/swcc file, use only CAA to start, stop, and restart SWCC.

The /usr/opt/SWCCx/scripts/swcc\_config utility includes options to enable, disable, and restart SWCC. These options call the /sbin/init.d/swcc script to start, stop, and restart SWCC. If you modified the /sbin/init.d/swcc script as described in Edit the Startup Script, these commands will no longer work on a cluster member.

---

**Note:** The recommended edit to the `/sbin/init.d/swcc` file prevents the `/usr/opt/SWCCx/scripts/swcc_config` utility from starting, stopping, or restarting SWCC when the `swcc` resource is registered with CAA. However, `swcc_config` is not aware of this change, and still attempts to start, stop, and restart SWCC if you choose these options. This action results in a problem-status message.

---

If you did not modify the `/sbin/init.d/swcc` script as described in Edit the Startup Script, do not use these options because they can interfere with CAA control of the `swcc` resource. For example, if you use `/usr/opt/SWCCx/scripts/swcc_config` to stop SWCC on a cluster member, CAA then restarts the `swcc` resource on that member or another member.

## Verifying Success

After you apply the procedure for using `swcc` in a cluster, you can verify whether it was successful by performing the following steps. You will probably find it easier to call the CAA commands interactively to test your work and get immediate feedback.

In this discussion, the active member is the member on which CAA starts the resource; it may or may not be the member that you are currently using.

To determine whether the CAA resource is online, and, if so, on which member, use the `caa_stat` command.

If the resource is not online, use the `caa_start` command to start the CAA resource. CAA reports the member on which the resource is started, as follows:

```
# caa_start swcc
Attempting to start `swcc` on member
`deli` Start of `swcc` on member `deli` succeeded.
```

To verify that `swcc` is running on the member, use the `ps` command to determine whether the `steamd` daemon is running on the member:

```
# ps agx | grep steamd
540567 pts/1 S 0:00.11 /usr/opt/SWCC5xx/bin/steamd sF
540850 pts/1 S + 0:00.00 grep steamd
```

To verify that the `swcc` agent is accessible from the client, use the client running on a PC to display information about the storage configuration. Make sure to use the cluster alias when connecting from the client.

To stop the CAA resource, use the `caa_stop swcc` command on any member.

To verify that swcc was removed, use the ps command to determine if the steamd daemon is running on the member. There should not be a steamd process.

```
# ps agx | grep steamd
540850 pts/1 S + 0:00.00 grep steamd
```

To start the CAA resource on another member, use the caa\_start swcc command with the -c option.

To verify that the swcc agent is accessible from the client, use the client running on a PC to display information about the storage configuration. Make sure to use the cluster alias when connecting from the client.

To fail over the CAA resource to another member of your choice, use the caa\_relocate swcc command with the -c option.

```
# caa_relocate swcc -c rye
Attempting to stop `swcc` on member `swiss`
Stop of `swcc` on member `swiss` succeeded.
Attempting to start `swcc` on member `rye`
Start of `swcc` on member `rye` succeeded.
# caa_stat swcc
NAME=swcc
TYPE=application
TARGET=ONLINE
STATE=ONLINE on rye
```

If you determine that the procedure was not successful, use the following table to identify and solve problems. In this discussion, the active member is the member on which you run the CAA resource.

**Table 22: Troubleshooting**

Problem	Possible Solutions
caa_stat command indicates that the CAA resource is offline	Use the caa_start swcc command interactively to start the CAA resource.
steamd daemon is running on multiple members	Use the caa_stat swcc to determine the cluster member on which CAA is running the resource. Then, use the ps command to determine the process IDs of the steamd daemon on the other members, and the kill command to stop these processes.

## glossary

This glossary defines terms pertaining to the ACS solution software. It is not a comprehensive glossary of computer terms

**8B/10B**

A type of byte definition encoding and decoding to reduce errors in data transmission patented by the IBM Corporation. This process of encoding and decoding data for transmission has been adopted by ANSI.

**adapter**

A device that converts the protocol and hardware interface of one bus type into another without changing the function of the bus.

**ACS**

*See* array controller software.

**AL\_PA**

*See* arbitrated loop physical address.

**alias address**

An AL\_PA value recognized by an arbitrated loop port in addition to the assigned AL\_PA.

**ANSI**

Pronounced “ann-see.” Acronym for the American National Standards Institute. An organization who develops standards used voluntarily by many manufacturers within the USA. ANSI is not a government agency.

**arbitrate**

A process of selecting one L\_Port from a collection of several ports that request use of the arbitrated loop concurrently.

**arbitrated loop**

A loop type of topology where two or more ports can be interconnected, but only two ports at a time can communicate.

**arbitrated loop physical address**

Abbreviated AL\_PA. A one-byte value used to identify a port in an Arbitrated Loop topology.

**array controller**

*See* controller.

**array controller software**

Abbreviated ACS. Software contained on a removable ROM program card that provides the operating system for the array controller.

<b>association set</b>	A group of remote copy sets that share selectable attributes for logging and failover. Members of an association set transition to the same state simultaneously. For example, if one association set member assumes the failsafe locked condition, then other members of the association set also assume the failsafe locked condition.
	An association set can also be used to share a log between a group of remote copy set members that require efficient use of the log space.
<b>asynchronous</b>	Pertaining to events that are scheduled as the result of a signal asking for the event; pertaining to that which is without any specified time relation. <i>See also</i> synchronous.
<b>autospare</b>	A controller feature that automatically replaces a failed disk drive. To aid the controller in automatically replacing failed disk drives, you can enable the AUTOSPARE switch for the failedset causing physically replaced disk drives to be automatically placed into the spareset. Also called “AUTONEWSPARE.”
<b>bad block</b>	A data block that contains a physical defect.
<b>bad block replacement</b>	Abbreviated BBR. A replacement routine that substitutes defect-free disk blocks for those found to have defects. This process takes place in the controller, transparent to the host.
<b>backplane</b>	The electronic printed circuit board into which you plug subsystem devices—for example, the SBB or power supply.
<b>battery hysteresis</b>	The ability of the software to allow write-block caching during the time a battery is charging, but only when a previous down time has not drained more than 50 percent of rated battery capacity.
<b>BBR</b>	<i>See</i> bad block replacement.
<b>BIST</b>	<i>See</i> built-in self-test.
<b>bit</b>	A single binary digit having a value of either 0 or 1. A bit is the smallest unit of data a computer can process.
<b>block</b>	Also called a sector. The smallest collection of consecutive bytes addressable on a disk drive. In integrated storage elements, a block contains 512 bytes of data, error codes, flags, and the block address header.
<b>bootstrapping</b>	A method used to bring a system or device into a defined state by means of its own action. For example, a machine routine whose first few instructions are enough to bring the rest of the routine into the computer from an input device.

---

<b>built-in self-test</b>	A diagnostic test performed by the array controller software on the controller policy processor.
<b>byte</b>	A binary character string made up of 8 bits operated on as a unit.
<b>cache memory</b>	A portion of memory used to accelerate read and write operations.
<b>cache module</b>	A fast storage buffer
<b>CCL</b>	CCL-Command Console LUN, a “SCSI Logical Unit Number” virtual-device used for communicating with Command Console Graphical User Interface (GUI) software.
<b>channel</b>	An interface that allows high speed transfer of large amounts of data. Another term for a SCSI bus. <i>See also</i> SCSI.
<b>chunk</b>	A block of data written by the host.
<b>chunk size</b>	The number of data blocks, assigned by a system administrator, written to the primary RAIDset or stripeset member before the remaining data blocks are written to the next RAIDset or stripeset member.
<b>CLCP</b>	An abbreviation for code-load code-patch utility. This utility is used to upgrade the controller and EMU software. It can also be used to patch the controller software.
<b>CLI</b>	See Command Line Interface.
<b>coax</b>	A two-conductor wire in which one conductor completely wraps the other with the two separated by insulation.
<b>cold swap</b>	A method of device replacement that requires the entire subsystem to be turned off before the device can be replaced. <i>See also</i> hot swap and warm swap.
<b>command line interface</b>	CLI. A command line entry utility used to interface with the HS-series controllers. CLI enables the configuration and monitoring of a storage subsystem through textual commands.
<b>concat commands</b>	Concat commands implement storageset expansion features.
<b>configuration file</b>	A file that contains a representation of a storage subsystem configuration.
<b>container</b>	1) Any entity that is capable of storing data, whether it is a physical device or a group of physical devices. (2) A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers the controller uses to create units.

<b>controller</b>	A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array. The HSG80 family controllers are examples of array controllers.
<b>copying</b>	A state in which data to be copied to the mirrorset is inconsistent with other members of the mirrorset. <i>See also</i> normalizing.
<b>copying member</b>	Any member that joins the mirrorset after the mirrorset is created is regarded as a copying member. Once all the data from the normal member (or members) is copied to a normalizing or copying member, the copying member then becomes a normal member. <i>See also</i> normalizing member.
<b>CSR</b>	An acronym for control and status register.
<b>daemon</b>	Pronounced “demon.” A program usually associated with a UNIX systems that performs a utility (housekeeping or maintenance) function without being requested or even known of by the user. A daemon is a diagnostic and execution monitor.
<b>data center cabinet</b>	A generic reference to large subsystem cabinets, such as the cabinets in which StorageWorks components can be mounted.
<b>data striping</b>	The process of segmenting logically sequential data, such as a single file, so that segments can be written to multiple physical devices (usually disk drives) in a round-robin fashion. This technique is useful if the processor is capable of reading or writing data faster than a single disk can supply or accept the data. While data is being transferred from the first disk, the second disk can locate the next segment.
<b>DDL</b>	Dual data link. The ability to operate on the CI bus using both paths simultaneously to the same remote node.
<b>device</b>	<i>See</i> node and peripheral device.
<b>differential I/O module</b>	A 16-bit I/O module with SCSI bus converter circuitry for extending a differential SCSI bus. <i>See also</i> I/O module.
<b>differential SCSI bus</b>	A bus in which a signal level is determined by the potential difference between two wires. A differential bus is more robust and less subject to electrical noise than is a single-ended bus.
<b>DIMM</b>	Dual inline Memory Module.
<b>dirty data</b>	The write-back cached data that has not been written to storage media, even though the host operation processing the data has completed.
<b>DMA</b>	Direct Memory Access.

---

<b>DOC</b>	DWZZA-On-a-Chip. ASCSI bus extender chip used to connect a SCSI bus in an expansion cabinet to the corresponding SCSI bus in another cabinet (See DWZZA).
<b>driver</b>	A hardware device or a program that controls or regulates another device. For example, a device driver is a driver developed for a specific device that allows a computer to operate with the device, such as a printer or a disk drive.
<b>dual-redundant configuration</b>	A controller configuration consisting of two active controllers operating as a single controller. If one controller fails, the other controller assumes control of the failing controller devices.
<b>dual-simplex</b>	A communications protocol that allows simultaneous transmission in both directions in a link, usually with no flow control.
<b>DUART</b>	Dual universal asynchronous receiver and transmitter. An integrated circuit containing two serial, asynchronous transceiver circuits.
<b>DWZZA</b>	A StorageWorks SCSI bus signal converter used to connect 8-bit single-ended devices to hosts with 16-bit differential SCSI adapters. This converter extends the range of a single-ended SCSI cable to the limit of a differential SCSI cable.
<b>DWZZB</b>	A StorageWorks SCSI bus signal converter used to connect a variety of 16-bit single-ended devices to hosts with 16-bit differential SCSI adapters.
<b>ECB</b>	External cache battery. The unit that supplies backup power to the cache module in the event the primary power source fails or is interrupted.
<b>ECC</b>	Error checking and correction.
<b>EDC</b>	Error detection code.
<b>EIA</b>	The abbreviation for Electronic Industries Association. EIA is a standards organization specializing in the electrical and functional characteristics of interface equipment.
<b>EMU</b>	Environmental monitoring unit. A unit that provides increased protection against catastrophic failures. Some subsystem enclosures include an EMU which works with the controller to detect conditions such as failed power supplies, failed blowers, elevated temperatures, and external air sense faults. The EMU also controls certain cabinet hardware including DOC chips, alarms, and fan speeds.

<b>ESD</b>	Electrostatic discharge. The discharge of potentially harmful static electrical voltage as a result of improper grounding.
<b>extended subsystem</b>	A subsystem in which two cabinets are connected to the primary cabinet.
<b>external cache battery</b>	<i>See ECB.</i>
<b>F_Port</b>	A port in a fabric where an N_Port or NL_Port may attach.
<b>fabric</b>	A group of interconnections between ports that includes a fabric element.
<b>failback</b>	The process of restoring data access to the newly-restored controller in a dual-redundant controller configuration. See also failover.
<b>failedset</b>	A group of failed mirrorset or RAIDset devices automatically created by the controller.
<b>failover</b>	The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.
	The ability for HSG80 controllers to transfer control from one controller to another in the event of a controller failure. This ensures uninterrupted operation. Use Transparent Failover mode for single HBA configurations. Use multiple-bus failover mode for Secure Path based configurations.
<b>FCA</b>	Fibre Channel Adapter
<b>FC-AL</b>	The Fibre Channel Arbitrated Loop standard. See <i>Fibre Channel</i> .
<b>FC-ATM</b>	ATM AAL5 over Fibre Channel
<b>FC-FG</b>	Fibre Channel Fabric Generic Requirements
<b>FG-FP</b>	Fibre Channel Framing Protocol (HIPPI on FC)
<b>FC-GS-1</b>	Fibre Channel Generic Services-1
<b>FC-GS-2</b>	Fibre Channel Generic Services-2
<b>FC-IG</b>	Fibre Channel Implementation Guide
<b>FC-IE</b>	Fibre Channel Link Encapsulation (ISO 8802.2)
<b>FC-PH</b>	The Fibre Channel Physical and Signaling standard.
<b>FC-SB</b>	Fibre Channel Single Byte Command Code Set
<b>FC-SW</b>	Fibre Channel Switched Topology and Switch Controls

---

<b>FCC</b>	Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.
<b>FCC Class A</b>	This certification label appears on electronic devices that can only be used in a commercial environment within the United States.
<b>FCC Class B</b>	This certification label appears on electronic devices that can be used in either a home or a commercial environment within the United States.
<b>FCP</b>	The mapping of SCSI-3 operations to Fibre Channel.
<b>FDDI</b>	Fiber Distributed Data Interface. An ANSI standard for 100 megabaud transmission over fiber optic cable.
<b>FD SCSI</b>	The fast, narrow, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s. <i>See also</i> FWD SCSI and SCSI.
<b>fiber</b>	A fiber or optical strand. Spelled <i>fibre</i> in Fibre Channel.
<b>fiber optic cable</b>	A transmission medium designed to transmit digital signals in the form of pulses of light. Fiber optic cable is noted for its properties of electrical isolation and resistance to electrostatic contamination.
<b>Fibre Channel</b>	A high speed, high-bandwidth serial protocol for channels and networks that interconnect over twisted pair wires, coaxial cable or fiber optic cable. The Fibre Channel Switched (FC-SW) (fabric) offers up to 16 million ports with cable lengths of up to 10 kilometers. The Fibre Channel Arbitrated Loop (FC-AL) topology offers speeds of up to 100 Mbytes/seconds and up to 127 nodes, all connected in serial. In contrast to SCSI technology, Fibre Channel does not require ID switches or terminators. The FC-AL loop may be connected to a Fibre Channel fabric for connection to other nodes.
<b>Fibre Channel topology</b>	An interconnection scheme that allows multiple Fibre Channel ports to communicate with each other. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.
<b>FL_Port</b>	A port in a fabric where N_Port or an NL_Port may be connected.
<b>flush</b>	The act of writing dirty data from cache to a storage media.
<b>FMU</b>	Fault management utility.
<b>forced errors</b>	A data bit indicating a corresponding logical data block contains unrecoverable data.
<b>frame</b>	An invisible unit used to transfer information in Fibre Channel.

<b>FRU</b>	Field replaceable unit. A hardware component that can be replaced at the customer location by service personnel or qualified customer service personnel.
<b>FRUTIL</b>	Field Replacement utility.
<b>full duplex (n)</b>	A communications system in which there is a capability for 2-way transmission and acceptance between two sites at the same time.
<b>full duplex (adj)</b>	Pertaining to a communications method in which data can be transmitted and received at the same time.
<b>FWD SCSI</b>	A fast, wide, differential SCSI bus with a maximum 16-bit data transfer rate of 20 MB/s. <i>See also</i> SCSI and FD SCSI.
<b>GBIC</b>	Gigabit Interface Converter. GBICs convert electrical signals to optical signals (and vice-versa.) They are inserted into the ports of the Fibre Channel switch and hold the Fibre Channel cables.
<b>GLM</b>	Gigabit link module
<b>giga</b>	A prefix indicating a billion ( $10^9$ ) units, as in gigabaud or gigabyte.
<b>gigabaud</b>	An encoded bit transmission rate of one billion ( $10^9$ ) bits per second.
<b>gigabyte</b>	A value normally associated with a disk drives storage capacity, meaning a billion ( $10^9$ ) bytes. The decimal value 1024 is usually used for one thousand.
<b>half-duplex (adj)</b>	Pertaining to a communications system in which data can be either transmitted or received but only in one direction at one time.
<b>hard address</b>	The AL_PA which an NL_Port attempts to acquire during loop initialization.
<b>heterogeneous host support</b>	Also called <i>noncooperating host support</i> .
<b>HIPPI-FC</b>	Fibre Channel over HIPPI
<b>host</b>	The primary or controlling computer to which a storage subsystem is attached.
<b>host adapter</b>	A device that connects a host system to a SCSI bus. The host adapter usually performs the lowest layers of the SCSI protocol. This function may be logically and physically integrated into the host system.
<b>HBA</b>	Host bus adapter

<b>host compatibility mode</b>	A setting used by the controller to provide optimal controller performance with specific operating systems. This improves the controller performance and compatibility with the specified operating system.
<b>hot disks</b>	A disk containing multiple hot spots. Hot disks occur when the workload is poorly distributed across storage devices which prevents optimum subsystem performance. <i>See also</i> hot spots.
<b>hot spots</b>	A portion of a disk drive frequently accessed by the host. Because the data being accessed is concentrated in one area, rather than spread across an array of disks providing parallel access, I/O performance is significantly reduced. <i>See also</i> hot disks.
<b>hot swap or hot-pluggable</b>	A method of device replacement that allows normal I/O activity on a device bus to remain active during device removal and insertion. The device being removed or inserted is the only device that cannot perform operations during this process. <i>See also</i> cold swap and warm swap.
<b>hub</b>	A device (concentrator) which performs some or all of the following functions: <ul style="list-style-type: none"><li>■ Automatic insertion of operational loop devices without disrupting the existing configuration.</li><li>■ Automatic removal of failed loop devices without impacting the existing configuration.</li><li>■ Provides a centralized (star) wiring configuration and maintenance point.</li><li>■ Provides central monitoring and management.</li></ul>
<b>IBR</b>	Initial Boot Record.
<b>ILF</b>	Illegal function.
<b>INIT</b>	Initialize input and output.
<b>initiator</b>	A SCSI device that requests an I/O process to be performed by another SCSI device, namely, the SCSI target. The controller is the initiator on the device bus. The host is the initiator on the host bus.
<b>instance code</b>	A four-byte value displayed in most text error messages and issued by the controller when a subsystem error occurs. The instance code indicates when during software processing the error was detected.
<b>interface</b>	A set of protocols used between components, such as cables, connectors, and signal levels.

<b>I/O</b>	Refers to input and output functions.
<b>I/O driver</b>	The set of code in the kernel that handles the physical I/O to a device. This is implemented as a fork process. Same as driver.
<b>I/O interface</b>	<i>See</i> interface.
<b>I/O module</b>	A 16-bit SBB shelf device that integrates the SBB shelf with either an 8-bit single ended, 16-bit single-ended, or 16-bit differential SCSI bus (see SBB).
<b>I/O operation</b>	The process of requesting a transfer of data from a peripheral device to memory (or visa versa), the actual transfer of the data, and the processing and overlaying activity to make both of those happen.
<b>IPI</b>	Intelligent Peripheral Interface. An ANSI standard for controlling peripheral devices by a host computer.
<b>IPI-3 Disk</b>	Intelligent Peripheral Interface Level 3 for Disk
<b>IPI-3 Tape</b>	Intelligent Peripheral Interface Level 3 for Tape
<b>JBOD</b>	Just a bunch of disks. A term used to describe a group of single-device logical units.
<b>kernel</b>	The most privileged processor access mode.
<b>LBN</b>	Logical Block Number.
<b>L_port</b>	A node or fabric port capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
<b>LED</b>	Light Emitting Diode.
<b>link</b>	A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.
<b>local connection</b>	A connection to the subsystem using either its serial maintenance port or the host SCSI bus. A local connection enables you to connect to one subsystem controller within the physical range of the serial or host SCSI cable.
<b>local terminal</b>	A terminal plugged into the EIA-423 maintenance port located on the front bezel of the controller. <i>See also</i> maintenance terminal.
<b>logical bus</b>	A single-ended bus connected to a differential bus by a SCSI bus signal converter.
<b>logical unit</b>	A physical or virtual device addressable through a target ID number. LUNs use their target bus connection to communicate on the SCSI bus.

---

<b>logical unit number</b>	LUN. A value that identifies a specific logical unit belonging to a SCSI target ID number. A number associated with a physical device unit during a task I/O operations. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.
<b>logon</b>	Also called login. A procedure whereby a participant, either a person or network connection, is identified as being an authorized network participant.
<b>loop</b>	<i>See</i> arbitrated loop.
<b>loop_ID</b>	A seven-bit value numbered contiguously from zero to 126-decimal and represent the 127 legal AL_PA values on a loop (not all of the 256 hex values are allowed as AL_PA values per FC-AL.)
<b>loop tenancy</b>	The period of time between the following events: when a port wins loop arbitration and when the port returns to a monitoring state.
<b>L_Port</b>	A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
<b>LUN</b>	<i>See</i> logical unit number.
<b>LRU</b>	Least recently used. A cache term used to describe the block replacement policy for read cache.
<b>Mbps</b>	Approximately one million ( $10^6$ ) bits per second—that is, megabits per second.
<b>maintenance terminal</b>	An EIA-423-compatible terminal used with the controller. This terminal is used to identify the controller, enable host paths, enter configuration information, and check the controller status. The maintenance terminal is not required for normal operations. <i>See also</i> local terminal.
<b>member</b>	A container that is a storage element in a RAID array.
<b>metadata</b>	The data written to a disk for the purposes of controller administration. Metadata improves error detection and media defect management for the disk drive. It is also used to support storage set configuration and partitioning. Nontransportable disks also contain metadata to indicate they are uniquely configured for StorageWorks environments. Metadata can be thought of as “data about data.”
<b>mirroring</b>	The act of creating an exact copy or image of data.

<b>mirrored write-back caching</b>	A method of caching data that maintains two copies of the cached data. The copy is available if either cache module fails.
<b>mirrorset</b>	<i>See</i> RAID level 1.
<b>MIST</b>	Module Integrity Self-Test.
<b>multibus failover</b>	Allows the host to control the failover process by moving the units from one controller to another.
<b>N_port</b>	A port attached to a node for use with point-to-point topology or fabric topology.
<b>NL_port</b>	A port attached to a node for use in all topologies.
<b>network</b>	In data communication, a configuration in which two or more terminals or devices are connected to enable information transfer.
<b>node</b>	In data communications, the point at which one or more functional units connect transmission lines.
<b>Non-L_Port</b>	A Node of Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports loop-capable ports.
<b>non-participating mode</b>	A mode within an L_Port that inhibits the port from participating in loop activities. L_Ports in this mode continue to retransmit received transmission words but are not permitted to arbitrate or originate frames. An L_Port in non-participating mode may or may not have an AL_PA. <i>See also</i> participating mode.
<b>nominal membership</b>	The desired number of mirrorset members when the mirrorset is fully populated with active devices. If a member is removed from a mirrorset, the actual number of members may fall below the “nominal” membership.
<b>node</b>	In data communications, the point at which one or more functional units connect transmission lines. In Fibre Channel, a device that has at least one N_Port or NL_Port.
<b>nonredundant controller configuration</b>	(1) A single controller configuration. (2) A controller configuration that does not include a second controller.
<b>normal member</b>	A mirrorset member that, block-for-block, contains the same data as other normal members within the mirrorset. Read requests from the host are always satisfied by normal members.

---

<b>normalizing</b>	Normalizing is a state in which, block-for-block, data written by the host to a mirrorset member is consistent with the data on other normal and normalizing members. The normalizing state exists only after a mirrorset is initialized. Therefore, no customer data is on the mirrorset.
<b>normalizing member</b>	A mirrorset member whose contents are the same as all other normal and normalizing members for data that has been written since the mirrorset was created or lost cache data was cleared. A normalizing member is created by a normal member when either all of the normal members fail or all of the normal members are removed from the mirrorset. <i>See also</i> copying member.
<b>NVM</b>	Non-Volatile Memory. A type of memory where the contents survive power loss. Also sometimes referred to as NVMEM.
<b>OCP</b>	Operator control panel. The control or indicator panel associated with a device. The OCP is usually mounted on the device and is accessible to the operator.
<b>offset</b>	A relative address referenced from the base element address. Event Sense Data Response Templates use offsets to identify various information contained within one byte of memory (bits 0 through 7).
<b>other controller</b>	The controller in a dual-redundant pair that is connected to the controller serving the current CLI session. <i>See also</i> this controller.
<b>outbound fiber</b>	One fiber in a link that carries information away from a port.
<b>parallel data transmission</b>	A data communication technique in which more than one code element (for example, bit) of each byte is sent or received simultaneously.
<b>parity</b>	A method of checking if binary numbers or characters are correct by counting the ONE bits. In odd parity, the total number of ONE bits must be odd; in even parity, the total number of ONE bits must be even.
<b>parity bit</b>	A binary digit added to a group of bits that checks to see if errors exist in the transmission.
<b>parity check</b>	A method of detecting errors when data is sent over a communications line. With even parity, the number of ones in a set of binary data should be even. With odd parity, the number of ones should be odd.
<b>parity RAID</b>	<i>See</i> RAIDset.
<b>participating mode</b>	A mode within an L_Port that allows the port to participate in loop activities. A port must have a valid AL_PA to be in participating mode.
<b>partition</b>	A logical division of a container, represented to the host as a logical unit.

<b>PCMCIA</b>	Personal Computer Memory Card Industry Association. An international association formed to promote a common standard for PC card-based peripherals to be plugged into notebook computers. The card commonly known as a PCMCIA card is about the size of a credit card.
<b>PDU</b>	Power distribution unit. The power entry device for StorageWorks cabinets. The CDU provides the connections necessary to distribute power to the cabinet shelves and fans.
<b>peripheral device</b>	Any unit, distinct from the CPU and physical memory, that can provide the system with input or accept any output from it. Terminals, printers, tape drives, and disks are peripheral devices.
<b>pluggable</b>	A replacement method that allows the complete system to remain online during device removal or insertion. The system bus must be halted, or quiesced, for a brief period of time during the replacement procedure. See also hot-pluggable.
<b>point-to-point connection</b>	A network configuration in which a connection is established between two, and only two, terminal installations. The connection may include switching facilities.
<b>port</b>	(1) In general terms, a logical channel in a communications system. (2) The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus.
	Regarding the controller, the port is (1) the logical route for data in and out of a controller that can contain one or more channels, all of which contain the same type of data. (2) The hardware and software that connects a controller to a SCSI device.
<b>port_name</b>	A 64-bit unique identifier assigned to each Fibre Channel port. The Port_Name is communicated during the login and port discovery process.
<b>preferred address</b>	The AL_PA which an NL_Port attempts to acquire first during initialization.
<b>primary cabinet</b>	The primary cabinet is the subsystem enclosure that contains the controllers, cache modules, external cache batteries, and the PVA module.
<b>private NL_Port</b>	An NL_Port which does not attempt login with the fabric and only communicates with NL_Ports on the same loop.
<b>program card</b>	The PCMCIA card containing the controller operating software.

<b>protocol</b>	The conventions or rules for the format and timing of messages sent and received.
<b>PTL</b>	Port-Target-LUN. The controller method of locating a device on the controller device bus.
<b>PVA module</b>	Power Verification and Addressing module.
<b>quiesce</b>	The act of rendering bus activity inactive or dormant. For example, “quiesce the SCSI bus operations during a device warm-swap.”
<b>RAID</b>	Redundant Array of Independent Disks. Represents multiple levels of storage access developed to improve performance or availability or both.
<b>RAID level 0</b>	A RAID storageset that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, enabling parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storagesets are sometimes referred to as stripesets.
<b>RAID level 0+1</b>	A RAID storageset that stripes data across an array of disks (RAID level 0) and mirrors the striped data (RAID level 1) to provide high I/O performance and high availability. This RAID level is alternatively called a striped mirrorset.
<b>RAID level 1</b>	A RAID storageset of two or more physical disks that maintain a complete and independent copy of the entire virtual disk's data. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are sometimes referred to as mirrorsets.
<b>RAID level 3</b>	A RAID storageset that transfers data parallel across the array disk drives a byte at a time, causing individual blocks of data to be spread over several disks serving as one enormous virtual disk. A separate redundant check disk for the entire array stores parity on a dedicated disk drive within the storageset. <i>See also</i> RAID level 5.
<b>RAID level 5</b>	A RAID storageset that, unlike RAID level 3, stores the parity information across all of the disk drives within the storageset. <i>See also</i> RAID level 3.

<b>RAID level 3/5</b>	A RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID.
<b>RAIDset</b>	<i>See</i> RAID level 3/5.
<b>RAM</b>	Random access memory.
<b>read ahead caching</b>	A caching technique for improving performance of synchronous sequential reads by prefetching data from disk.
<b>read caching</b>	A cache management method used to decrease the subsystem response time to a read request by allowing the controller to satisfy the request from the cache memory rather than from the disk drives.
<b>reconstruction</b>	The process of regenerating the contents of a failed member data. The reconstruct process writes the data to a spareset disk and incorporates the spareset disk into the mirrorset, striped mirrorset, or RAIDset from which the failed member came. <i>See also</i> regeneration.
<b>reduced</b>	Indicates that a mirrorset or RAIDset is missing one member because the member has failed or has been physically removed.
<b>redundancy</b>	The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors. A RAIDset is considered to be redundant when user data is recorded directly to one member and all of the other members include associated parity information.
<b>regeneration</b>	(1) The process of calculating missing data from redundant data. (2) The process of recreating a portion of the data from a failing or failed drive using the data and parity information from the other members within the storageset. The regeneration of an entire RAIDset member is called reconstruction. <i>See also</i> reconstruction.
<b>remote copy</b>	A feature intended for disaster tolerance and replication of data from one storage subsystem or physical site to another subsystem or site. Remote copy also provides methods of performing a backup at either the local or remote site. With remote copy, user applications continue to run while data movement goes on in the background. Data warehousing, continuous computing, and enterprise applications all require remote copy capabilities.

---

<b>remote copy set</b>	A bound set of two units, one located locally and one located remotely, for long-distance mirroring. The units can be a single disk, or a storageset, mirrorset, or RAIDset. A unit on the local controller is designated as the “initiator” and a corresponding unit on the remote controller is designated as the “target”.
<b>request rate</b>	The rate at which requests are arriving at a servicing entity.
<b>RFI</b>	Radio frequency interference. The disturbance of a signal by an unwanted radio signal or frequency.
<b>replacement policy</b>	The policy specified by a switch with the SET FAILEDSET command indicating whether a failed disk from a mirrorset or RAIDset is to be automatically replaced with a disk from the spareset. The two switch choices are AUTOSPARE and NOAUTOSPARE.
<b>SBB</b>	StorageWorks building block. (1) A modular carrier plus the interface required to mount the carrier into a standard StorageWorks shelf. (2) any device conforming to shelf mechanical and electrical standards installed in a 3.5-inch or 5.25-inch carrier, whether it is a storage device or power supply.
<b>SCSI</b>	Small computer system interface. (1) An ANSI interface standard defining the physical and electrical parameters of a parallel I/O bus used to connect initiators to devices. (2) a processor-independent standard protocol for system-level interfacing between a computer and intelligent devices including hard drives, floppy disks, CD-ROMs, printers, scanners, and others.
<b>SCSI-A cable</b>	A 50-conductor (25 twisted-pair) cable generally used for single-ended, SCSI-bus connections.
<b>SCSI bus signal converter</b>	Sometimes referred to as an adapter. (1) A device used to interface between the subsystem and a peripheral device unable to be mounted directly into the SBB shelf of the subsystem. (2) a device used to connect a differential SCSI bus to a single-ended SCSI bus. (3) A device used to extend the length of a differential or single-ended SCSI bus. <i>See also</i> DOC (DWZZA-On-a-chip) and I/O module.
<b>SCSI device</b>	(1) A host computer adapter, a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus. (2) Any physical unit that can communicate on a SCSI bus.
<b>SCSI device ID number</b>	A bit-significant representation of the SCSI address referring to one of the signal lines, numbered 0 through 7 for an 8-bit bus, or 0 through 15 for a 16-bit bus. <i>See also</i> target ID number.

<b>SCSI ID number</b>	The representation of the SCSI address that refers to one of the signal lines numbered 0 through 15.
<b>SCSI-P cable</b>	A 68-conductor (34 twisted-pair) cable generally used for differential bus connections.
<b>SCSI port</b>	(1) Software: The channel controlling communications to and from a specific SCSI bus in the system. (2) Hardware: The name of the logical socket at the back of the system unit to which a SCSI device is connected.
<b>serial transmission</b>	A method transmission in which each bit of information is sent sequentially on a single channel rather than simultaneously as in parallel transmission.
<b>signal converter</b>	<i>See</i> SCSI bus signal converter.
<b>single ended I/O module</b>	A 16-bit I/O module. <i>See also</i> I/O module.
<b>single-ended SCSI bus</b>	An electrical connection where one wire carries the signal and another wire or shield is connected to electrical ground. Each signal logic level is determined by the voltage of a single wire in relation to ground. This is in contrast to a differential connection where the second wire carries an inverted signal.
<b>spareset</b>	A collection of disk drives made ready by the controller to replace failed members of a storage set.
<b>storage array</b>	An integrated set of storage devices.
<b>storage array subsystem</b>	<i>See</i> storage subsystem.
<b>storage set</b>	(1) A group of devices configured with RAID techniques to operate as a single container. (2) Any collection of containers, such as stripesets, mirrorsets, striped mirrorsets, and RAIDsets.
<b>storage set expansion</b>	The dynamic expansion of the storage capacity (size) of a unit. A storage container is created in the form of a concatenation set which is added to the existing storage set defined as a unit.
<b>storage subsystem</b>	The controllers, storage devices, shelves, cables, and power supplies used to form a mass storage subsystem.
<b>storage unit</b>	The general term that refers to storage sets, single-disk units, and all other storage devices that are installed in your subsystem and accessed by the host. A storage unit can be any entity that is capable of storing data, whether it is a physical device or a group of physical devices.

---

<b>StorageWorks</b>	A family of modular data storage products that allow customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in StorageWorks enclosures to form storage subsystems.
	StorageWorks systems include integrated SBBs and array controllers to form storage subsystems. System-level enclosures to house the shelves and standard mounting devices for SBBs are also included.
<b>stripe</b>	The data divided into blocks and written across two or more member disks in an array.
<b>striped mirrorset</b>	<i>See</i> RAID level 0+1.
<b>stripeset</b>	<i>See</i> RAID level 0.
<b>stripe size</b>	The stripe capacity as determined by $n-1$ times the chunksize, where $n$ is the number of RAIDset members.
<b>striping</b>	The technique used to divide data into segments, also called chunks. The segments are striped, or distributed, across members of the stripeset. This technique helps to distribute hot spots across the array of physical devices to prevent hot spots and hot disks.
	Each stripeset member receives an equal share of the I/O request load, improving performance.
<b>surviving controller</b>	The controller in a dual-redundant configuration pair that serves its companion devices when the companion controller fails.
<b>switch</b>	A method that controls the flow of functions and operations in software.
<b>synchronous</b>	Pertaining to a method of data transmission which allows each event to operate in relation to a timing signal. <i>See also</i> asynchronous.
<b>tape</b>	A storage device supporting sequential access to variable sized data records.
<b>target</b>	(1) A SCSI device that performs an operation requested by an initiator. (2) Designates the target identification (ID) number of the device.
<b>target ID number</b>	The address a bus initiator uses to connect with a bus target. Each bus target is assigned a unique target address.
<b>this controller</b>	The controller that is serving your current CLI session through a local or remote terminal. <i>See also</i> other controller.

<b>tape inline exerciser</b>	(TILX) The controller diagnostic software to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.
<b>topology</b>	An interconnection scheme that allows multiple Fibre Channel ports to communicate with each other. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.
<b>transfer data rate</b>	The speed at which data may be exchanged with the central processor, expressed in thousands of bytes per second (kbytes).
<b>transparent failover</b>	Keeps the storage array available to the hosts by allowing the surviving controller of a dual redundant pair to take over total control of the subsystem and is transparent (invisible) to the hosts.
<b>ULP</b>	Upper Layer Protocol.
<b>ULP process</b>	A function executing within a Fibre Channel node which conforms to the Upper Layer Protocol (ULP) requirements when interacting with other ULP processes.
<b>Ultra SCSI</b>	A Fast-20 SCSI bus. <i>See also</i> Wide Ultra SCSI.
<b>unit</b>	A container made accessible to a host. A unit may be created from a single disk drive or tape drive. A unit may also be created from a more complex container such as a RAIDset. The controller supports a maximum of eight units on each target. <i>See also</i> target and target ID number.
<b>unwritten cached data</b>	Sometimes called unflushed data. <i>See</i> dirty data.
<b>UPS</b>	Uninterruptible power supply. A battery-powered power supply guaranteed to provide power to an electrical device in the event of an unexpected interruption to the primary power supply. Uninterruptible power supplies are usually rated by the amount of voltage supplied and the length of time the voltage is supplied.
<b>VHDCI</b>	Very high-density-cable interface. A 68-pin interface. Required for Ultra-SCSI connections.
<b>virtual terminal</b>	A software path from an operator terminal on the host to the controller's CLI interface, sometimes called a host console. The path can be established via the host port on the controller or via the maintenance port through an intermediary host.
<b>VTDPY</b>	An abbreviation for Virtual Terminal Display Utility.

<b>warm swap</b>	A device replacement method that allows the complete system to remain online during device removal or insertion. The system bus may be halted, or quiesced, for a brief period of time during the warm-swap procedure.
<b>Wide Ultra SCSI</b>	Fast/20 on a Wide SCSI bus.
<b>Worldwide name</b>	A unique 64-bit number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by manufacturing prior to shipping. This name is referred to as the node ID within the CLI.
<b>write-back caching</b>	A cache management method used to decrease the subsystem response time to write requests by allowing the controller to declare the write operation “complete” as soon as the data reaches its cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.
<b>write-through caching</b>	Write-through caching always writes directly to disk, ensuring that the application is never tricked into believing that the data is on the disk when it may not be. This results in highest data integrity, though with slightly reduced performance.
<b>write hole</b>	The period of time in a RAID level 1 or RAID level 5 write operation when an opportunity emerges for undetectable RAIDset data corruption. Write holes occur under conditions such as power outages, where the writing of multiple members can be abruptly interrupted. A battery backed-up cache design eliminates the write hole because data is preserved in cache and unsuccessful write operations can be retried.
<b>write-through cache</b>	A cache management technique for retaining host write requests in read cache. When the host requests a write operation, the controller writes data directly to the storage device. This technique allows the controller to complete some read requests from the cache, greatly improving the response time to retrieve data. The operation is complete only after the data to be written is received by the target storage device.  This cache management method may update, invalidate, or delete data from the cache memory accordingly, to ensure that the cache contains the most current data.



# index

## A

ADD CONNECTIONS  
    multiple-bus failover 46  
    transparent failover 44  
ADD UNIT  
    multiple-bus failover 46  
    transparent failover 44  
adding  
    virtual disks 224  
adding a disk drive to the spareset  
    configuration options 171  
adding disk drives  
    configuration options 171  
Agent  
    functions 132  
    installing 137  
array of disk drives 74  
assigning unit numbers 44  
assignment  
    unit numbers  
        fabric topology 169  
    unit qualifiers  
        fabric topology 169  
assignment of unit numbers  
    fabric topology  
        partition 169  
        single disk 169  
asynchronous event service 224  
autospare  
    enabling  
        fabric topology 172  
availability 81

## B

Back up, Clone, Move Data 187  
backup  
    cloning data 189  
    subsystem configuration 188

## C

cabling  
    controller pair 157  
    multiple-bus failover  
        fabric topology configuration 156  
    single controller 150  
cache modules  
    location 28, 29  
    read caching 36  
    write-back caching 36  
    write-through caching 36  
caching techniques  
    mirrored 37  
    read caching 36  
    read-ahead caching 36  
    write-back caching 36  
    write-through caching 36  
changing switches  
    configuration options 173, 174  
chunk size  
    choosing for RAIDsets and stripesets 89  
    controlling stripesize 89  
    using to increase request rate 90  
    using to increase write performance 91  
CHUNKSIZE 89  
CLI commands

installation verification 155, 162  
CLI configuration  
example 181  
CLI prompt  
changing  
  fabric topology 170  
Client  
  removing 223  
  uninstalling 223  
CLONE utility  
  backup 189  
cloning  
  backup 189  
command console LUN 38  
  SCSI-2 mode 47  
  SCSI-3 mode 47  
configuration  
  backup 188  
  fabric topology  
    devices 163  
    multiple-bus failover  
      cabling 156  
    multiple-bus failover using CLI 181  
    single controller cabling 149  
  restoring 91  
  rules 61  
Configuration Flowchart 22  
configuration options  
  fabric topology  
    adding a disk drive to the spareset 171  
    adding disk drives 171  
    changing switches  
      device 174  
      displaying the current switches 173  
      initialize 174  
      RAIDset and mirrorset 173  
      unit 174  
    changing switches for a storageset or  
      device 173  
    changing the CLI prompt 170  
    deleting a storageset 173  
    enabling autospare 172  
  removing a disk drive from the spareset  
    171  
  configuring  
    pager notification 224  
  configuring devices  
    fabric topology 163  
  connections 39  
    naming 39  
  containers  
    attributes 72  
    illustrated 73, 164  
    mirrorsets 80  
    planning storage 72  
    stripesets 78  
  controller  
    verification of installation 162  
  controller verification  
    installation 155, 162  
  controllers  
    cabling 150, 157  
    location 28, 29  
    node IDs 55  
    verification of installation 155, 162  
    worldwide names 55  
conventions  
  document 18  
creating  
  storageset and device profiles 75  
Creating Clones for Backup 189

**D**

DECsafe 112  
Destroy/Nodesstroy parameters 92  
device switches  
  changing  
    fabric topology 174  
devices  
  changing switches  
    fabric topology 173  
  configuration  
    fabric topology 163  
  creating a profile 75

disk drives  
adding  
  fabric topology 171  
adding to the spareset  
  fabric topology 171  
array 74  
  corresponding storagesets 93  
  dividing 85  
  removing from the spareset  
    fabric topology 171  
displaying the current switches  
  fabric topology 173  
dividing storagesets 85  
document  
  conventions 18

## E

enabling  
  switches 87  
erasing metadata 92  
establishing a local connection 148

## F

fabric topology  
  configuration  
    single controller cabling 149  
failover 31  
  multiple-bus 33  
  transparent 31  
First enclosure of multiple-enclosure subsystem  
  storage map template 1 200, 203, 205,  
    207, 209, 210, 212  
functions  
  Agent 132

## G

genvmunix, using 112  
geometry  
  initialize switches 93  
Geometry parameters 93

## H

Host access  
  restricting in multiple-bus failover mode  
    disabling access paths 52  
host access  
  restricting by offsets  
    multiple-bus failover 54  
    transparent failover 51  
  restricting in multiple-bus failover mode 52  
  restricting in transparent failover mode 49  
    disabling access paths 50  
    separate links 49  
host adapter  
  installation 104  
  preparation 104  
host connections 39  
  naming 39  
Host storage configuration  
  verify 175  
HSG Agent  
  install and configure 131  
  network connection 134  
  overview 133  
  reconfigure agent 143  
  remove agent 145

## I

initialize switches  
  changing  
    fabric topology 174  
CHUNKSIZE 89  
geometry 93  
NOSAVE\_CONFIGURATION 91  
SAVE\_CONFIGURATION 91  
Insight Manager 224  
installation  
  controller verification 155, 162  
  invalid network port assignments 219  
  there is no disk in the drive message 220  
installation verification  
  CLI commands 155, 162

installing  
  Agent 137  
integrating SWCC 224  
invalid network port assignments 219

## J

JBOD 74

## L

LOCATE  
  find devices 94  
location  
  cache module 28, 29  
  controller 28, 29  
LUN IDs  
  general description 57  
LUN presentation 45

## M

maintenance port connection  
  establishing a local connection 148  
  illustrated 148  
mapping storage sets 93  
messages  
  there is no disk in the drive 220  
mirrored caching  
  enabling 37  
  illustrated 37  
mirrorset switches  
  changing  
    fabric topology 173  
mirrorsets  
  planning considerations 80  
    important points 81  
  switches 88  
Model 2200 Storage Maps  
  examples 65  
moving storage sets 193  
multiple-bus failover  
  restricting host access by offsets 54  
multiple-bus failover 33

ADD CONNECTIONS command 46  
ADD UNIT command 46  
ADD UNITcommand 46  
CLI configuration procedure  
  fabric topology 181  
fabric topology  
  preferring units 170  
fabric topology configuration  
  cabling 156  
host connections 46  
restricting host access 52  
  disabling access paths 52  
SET CONNECTIONS command 46  
SET UNITcommand 46

## N

network port assignments 219  
new features 115  
node IDs 55  
restoring 56  
NODE\_ID  
  worldwide name 55  
NOSAVE\_CONFIGURATION 91

## O

offset  
  LUN presentation 45  
  restricting host access  
    multiple-bus failover 54  
    transparent failover 51  
  SCSI version factor 46  
online help  
  SWCC 224  
options  
  for mirrorsets 88  
  for RAIDsets 88  
  initialize 89  
  other controller 29

## P

pager notification 224

- configuring 224
- partitions
  - assigning a unit number
    - fabric topology 169
  - defining 86
  - planning considerations 85
    - guidelines 86
- performance 81
- Physical connection, making 104
- planning 59
  - overview 75
  - striped mirrorsets 84
  - stripesets 78
- planning configurations
  - where to start 60
- planning considerations 81
- planning storage
  - containers 72
- planning storagesets
  - characteristics
    - changing switches 87
    - enabling switches 87
    - initialization switch 87
    - storageset switch 87
    - unit switch 87
  - switches
    - initialization 89
    - storageset 88
- preferring units
  - multiple-bus failover
    - fabric topology 170
- profiles
  - creating 75
  - description 75
  - storageset 197
    - example 198
- R**
  - RAIDset switches
    - changing
      - fabric topology 173
  - RAIDsets
- choosing chunk size 89
- maximum membership 83
- planning considerations 81
  - important points 82
- switches 88
- read caching
  - enabled for all storage units 36
  - general description 36
- read requests
  - decreasing the subsystem response time with
    - read caching 36
- read-ahead caching 36
  - enabled for all disk units 36
- removing
  - Client 223
- request rate 90
- requirements
  - host adapter installation 104
- restricting host access
  - disabling access paths
    - multiple-bus failover 52
    - transparent failover 50
  - multiple-bus failover 52
- separate links
  - transparent failover 49
  - transparent failover 49
- S**
  - SAVE\_CONFIGURATION 91
  - saving configuration 91
- SCSI version
  - offset 46
- SCSI-2
  - assigning unit numbers 47
  - command console lun 47
- SCSI-3
  - assigning unit numbers 47
  - command console lun 47
- Second enclosure of multiple-enclosure subsystem
  - storage map template 2 201
- selective storage presentation 48

SET CONNECTIONS  
  multiple-bus failover [46](#)  
  transparent failover [44](#)

SET UNIT  
  multiple-bus failover [46](#)

setting  
  controller configuration handling [91](#)

single disk (JBOD)  
  assigning a unit number  
  fabric topology [169](#)

Single-enclosure subsystem  
  storage map template 1  
    [200](#)

storage  
  creating map [93](#)  
  profile  
    example [198](#)  
  storage configurations [59](#)  
  storage map [93](#)

Storage map template 1 [200](#)  
  first enclosure of multiple-enclosure  
    subsystem [200](#)  
  single enclosure subsystem [200](#)

Storage map template 2 [201](#)  
  second enclosure of multiple-enclosure  
    subsystem [201](#)

Storage map template 3 [202](#)  
  third enclosure of multiple-enclosure  
    subsystem [202](#)

Storage map template 4  
  first enclosure of multiple-enclosure  
    subsystem [203](#)

Storage map template 5  
  first enclosure of multiple-enclosure  
    subsystem [205](#)

Storage map template 6  
  first enclosure of multiple-enclosure  
    subsystem [207](#)

Storage map template 7  
  first enclosure of multiple-enclosure  
    subsystem [209](#)

Storage map template 8

first enclosure of multiple-enclosure  
  subsystem [210](#)

Storage map template 9  
  first enclosure of multiple-enclosure  
    subsystem [212](#)

storageset  
  deleting  
    fabric topology [173](#)  
  fabric topology  
    changing switches [173](#)  
  planning considerations [76](#)  
    mirrorsets [80](#)  
    partitions [85](#)  
    RAIDsets [81](#)  
    striped mirrorsets [83](#)  
    stripesets [76](#)  
  profile [75](#)  
  profiles [197](#)

storageset profile [75](#)

storageset switches  
  SET command [88](#)

storagesets  
  creating a profile [75](#)  
  moving [193](#)

striped mirrorsets  
  planning [84](#)  
  planning considerations [83](#)

stripesets  
  distributing members across buses [79](#)  
  planning [78](#)  
  planning considerations [76](#)  
    important points [78](#)

subsystem  
  saving configuration [91](#)

subsystem configuration  
  backup [188](#)

SWCC [132](#)  
  additional information [224](#)  
  integrating [224](#)  
  online help [224](#)

SWCC Agent  
  set up [225](#)

switches  
  changing 87  
  changing characteristics 87  
  CHUNKSIZE 89  
  enabling 87  
  mirrorsets 88  
  NOSAVE\_CONFIGURATION 91  
  RAIDset 88  
  SAVE\_CONFIGURATION 91  
switches for storagesets  
  overview 87

## T

templates  
  subsystem profile 197

terminology  
  other controller 29  
  this controller 29

Third enclosure of multiple-enclosure subsystem  
  storage map template 3 202

this controller 29

transparent failover 31

  ADD CONNECTIONS 44  
  ADD UNIT 44  
  matching units to host connections 44  
  restricting host access 49  
    disabling access paths 50  
    separate links 49  
  restricting host access by offsets 51  
  SET CONNECTIONS 44

troubleshooting  
  invalid network port assignments 219  
  there is no disk in the drive message 220

Tru64 UNIX utilities 112

## U

uninstalling  
  Client 223

unit numbers  
  assigning 44

  fabric topology 169  
  assigning depending on SCSI version 47  
  assigning in fabric topology  
    partition 169  
    single disk 169

unit qualifiers  
  assigning  
    fabric topology 169

unit switches  
  changing  
    fabric topology 174

units  
  LUN IDs 57

Upgrade procedures  
  solution software 113

## V

verification  
  controller installation 155, 162

verification of installation  
  controller 155, 162

virtual disks  
  adding 224

## W

worldwide names 55  
  NODE\_ID 55  
  REPORTED PORT\_ID 55  
  restoring 56

write performance 91

write requests  
  improving the subsystem response time with  
    write-back caching 36  
  placing data with with write-through caching  
    37

write-back caching  
  general description 36

write-through caching  
  general description 36